



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

A company has an internal web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group in a single Availability Zone. A SysOps administrator must make the application highly available. Which action should the SysOps administrator take to meet this requirement?

- A. Increase the maximum number of instances in the Auto Scaling group to meet the capacity that is required at peak usage.
- B. Increase the minimum number of instances in the Auto Scaling group to meet the capacity that is required at peak usage.
- C. Update the Auto Scaling group to launch new instances in a second Availability Zone in the same AWS Region.
- D. Update the Auto Scaling group to launch new instances in an Availability Zone in a second AWS Region.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Web_AmazonExams** Highly Voted 👍 11 months, 2 weeks ago

Selected Answer: C

Target tracking is only for unpredictable workloads. Answer is C.
upvoted 30 times

🗳️ 👤 **Finger41** Highly Voted 👍 3 years, 8 months ago

Option C is the answer. Spreading it across zones.

A and B doesn't make it highly available if a zone goes down.

D is false as ASG is region bound, as seen here: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/auto-scaling-benefits.html>
upvoted 30 times

🗳️ 👤 **Robertwilliamm** Most Recent 🕒 4 days, 4 hours ago

Selected Answer: C

C. is Correct Option
Thanks to SkillCertExams I successfully cleared my SOA-C02 exam today.
upvoted 1 times

🗳️ 👤 **renatom1976** 1 month, 2 weeks ago

Selected Answer: C

Option C is the correct option. Is not possible to stay in another region.
D is false.
upvoted 1 times

🗳️ 👤 **nerea_mineeys** 1 month, 3 weeks ago

Selected Answer: C

A and B discarded, C is the option since ASG is in a region. <https://ln.run/RMBh5>
upvoted 1 times

🗳️ 👤 **Xavier_Examen** 2 months, 3 weeks ago

Selected Answer: C

Answer is C.
upvoted 1 times

🗳️ 👤 **Molade** 3 months ago

Selected Answer: C

High availability
upvoted 1 times

🗳️ 👤 **heavyp** 3 months, 1 week ago

Selected Answer: C

A and B discarded, C is the option since ASG is in a region.

upvoted 1 times

🗨️ 👤 **ogum** 4 months, 3 weeks ago

Selected Answer: C

Option C

upvoted 1 times

🗨️ 👤 **numark** 5 months ago

Selected Answer: C

There are TWO SYSOPS exam question sets. This one which has 478 and the other which has 932. I am of course using this one because I THINK it is the most recent. Why is there two? If one is old why didn't you get rid of it yet? If they are both getting updated then why don't you combine them?

This sucks not knowing which one is the most recent and it should be resolved by exam tops ASAP!

upvoted 2 times

🗨️ 👤 **mk1523** 7 months ago

Selected Answer: C

Option C

upvoted 1 times

🗨️ 👤 **64rl0** 9 months, 3 weeks ago

Selected Answer: C

Answer is C.

upvoted 1 times

🗨️ 👤 **nicker003** 1 year, 2 months ago

Just took the exam yesterday and the pass (April-24) had about 10 to 15 questions that are not in this dump. But don't worry too much, study carefully and watch everyone's discussions with the answer has the most voted. You will have more knowledge and tips to do questions that are not in the dump. Wish everyone success and thank you Examtopics.

upvoted 6 times

🗨️ 👤 **Yowie351** 1 year, 4 months ago

Selected Answer: C

Option C

upvoted 2 times

🗨️ 👤 **tamng** 1 year, 6 months ago

C. Update the Auto Scaling group to launch new instances in a second Availability Zone in the same AWS Region.

upvoted 1 times

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

Selected Answer: C

Option C

upvoted 1 times

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

C. Update the Auto Scaling group to launch new instances in a second Availability Zone in the same AWS Region. Most Voted

upvoted 1 times

A company hosts a website on multiple Amazon EC2 instances that run in an Auto Scaling group. Users are reporting slow responses during peak times between 6 PM and 11 PM every weekend. A SysOps administrator must implement a solution to improve performance during these peak times. What is the MOST operationally efficient solution that meets these requirements?

- A. Create a scheduled Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to increase the desired capacity before peak times.
- B. Configure a scheduled scaling action with a recurrence option to change the desired capacity before and after peak times.
- C. Create a target tracking scaling policy to add more instances when memory utilization is above 70%.
- D. Configure the cooldown period for the Auto Scaling group to modify desired capacity before and after peak times.

Suggested Answer: B

Community vote distribution

B (100%)

 **Finger41** Highly Voted 3 years, 8 months ago

Answer is B - "Scheduled scaling helps you to set up your own scaling schedule according to predictable load changes. For example, let's say that every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can configure a schedule for Amazon EC2 Auto Scaling to increase capacity on Wednesday and decrease capacity on Friday."

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

upvoted 23 times

 **renatom1976** Most Recent 1 month, 2 weeks ago

Selected Answer: B

B. is the correct answer. Schedule time is the tool option that should be used in this situation.


upvoted 1 times

 **Xavier_Examen** 2 months, 3 weeks ago

Selected Answer: B

Answer is B

upvoted 1 times

 **heavyp** 3 months, 1 week ago

Selected Answer: B

Scheduling with the peak time, and defining ASG

upvoted 1 times

 **ogum** 4 months, 3 weeks ago

Selected Answer: B

B. Configure a scheduled scaling action with a recurrence option to change the desired capacity before and after peak times.

upvoted 1 times

 **64rl0** 9 months, 3 weeks ago

Selected Answer: B

Answer is B

upvoted 1 times

 **ft_cloud** 1 year, 8 months ago

Answer is B

upvoted 1 times

 **NAVADIYA** 1 year, 8 months ago

B. Configure a scheduled scaling action with a recurrence option to change the desired capacity before and after peak times.

upvoted 1 times

 **alexiscloud** 1 year, 8 months ago

Target tracking is only for unpredictable workloads. Answer is B.

upvoted 4 times

🗳️ 👤 **mimahmed_awseducate** 1 year, 7 months ago

Thanks

upvoted 1 times

🗳️ 👤 **habros** 1 year, 10 months ago

Selected Answer: B

I know when I want to scale up - predictable. Target tracking is only for unpredictable workloads. Lambda + EB is not the most op efficient. Cooldown does not apply here.

upvoted 1 times

🗳️ 👤 **DZRomero** 1 year, 11 months ago

Correcto; la respuesta es opción B

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbbbbbb

upvoted 1 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

upvoted 1 times

🗳️ 👤 **Mecdrex** 3 years, 2 months ago

Selected Answer: B

Vote B

upvoted 1 times

🗳️ 👤 **roka_ua** 3 years, 3 months ago

Selected Answer: B

Vote B

upvoted 2 times

🗳️ 👤 **MrkJobs** 3 years, 5 months ago

Selected Answer: B

as finger41 say

upvoted 1 times

🗳️ 👤 **szl0144** 3 years, 5 months ago

vote B

upvoted 1 times

A company is running a website on Amazon EC2 instances behind an Application Load Balancer (ALB). The company configured an Amazon CloudFront distribution and set the ALB as the origin. The company created an Amazon Route 53 CNAME record to send all traffic through the CloudFront distribution. As an unintended side effect, mobile users are now being served the desktop version of the website. Which action should a SysOps administrator take to resolve this issue?

- A. Configure the CloudFront distribution behavior to forward the User-Agent header.
- B. Configure the CloudFront distribution origin settings. Add a User-Agent header to the list of origin custom headers.
- C. Enable IPv6 on the ALB. Update the CloudFront distribution origin settings to use the dualstack endpoint.
- D. Enable IPv6 on the CloudFront distribution. Update the Route 53 record to use the dualstack endpoint.

Suggested Answer: C

Reference:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html>

Value/Route traffic to

Choose **Alias to Application and Classic Load Balancer** or **Alias to Network Load Balancer**, then choose the Region that the endpoint is from.

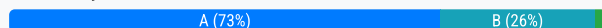
If you created the hosted zone and the ELB load balancer using the same AWS account –
Choose the name that you assigned to the load balancer when you created it.

If you created the hosted zone and the ELB load balancer using different accounts – Enter the value that you got in step 1 of this procedure.

Note

The console prepends **dualstack.** to the DNS name of the Application and Classic Load Balancer from the same AWS account only. When a client, such as a web browser, requests the IP address for your domain name (example.com) or subdomain name (www.example.com), the client can request an IPv4 address (an A record), an IPv6 address (a AAAA record), or both IPv4 and IPv6 addresses (in separate requests with IPv4 first). The **dualstack.** designation allows Route 53 to respond with the appropriate IP address for your load balancer based on which IP address format the client requested. You will need to prepend **dualstack.** for Application and Classic Load Balancer from the different account.

Community vote distribution



Pepepep Highly Voted 3 years ago

Selected Answer: A

I agree it is A because:

1. B is wrong, since you are modifying origin custom headers that are values that you set unilaterally, independent of the Header of the request that you received from client. As the documentation states, the uses cases for for origin custom headers are:

Identifying requests from CloudFront

Determining which requests come from a particular distribution

Enabling cross-origin resource sharing (CORS)

Controlling access to content

(<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/add-origin-custom-headers.html>)

upvoted 20 times

Pepepep 9 months, 1 week ago

2. A is correct. If you want CloudFront to cache different versions of your objects based on the device that a user is using to view your content, we recommend that you configure CloudFront to forward one or more of the following headers to your custom origin:

CloudFront-Is-Desktop-Viewer

CloudFront-Is-Mobile-Viewer

CloudFront-Is-SmartTV-Viewer

CloudFront-Is-Tablet-Viewer

As you can see based on what comes in the request received by CF, It is going to set the value to true before forwarding the request to your origin. The resume, with origin custom headers you are the one that decide what is going to be on the header, there is no way to match what it comes on the request received on CF and what is going to be sent to the origin. On the other hand with User-Agent header, CF inspects the header and determines what type of device is used (Smart TV, Tablet, Desktop, Mobile).

upvoted 16 times

  **dontcomplain** Highly Voted 3 years, 2 months ago

Why would this not be B? It would serve the same function more efficiently?



upvoted 10 times

  **Xavier_Examen** Most Recent 2 months, 3 weeks ago

Selected Answer: A

A is correct.



upvoted 1 times

  **ogum** 4 months, 3 weeks ago

Selected Answer: A

A. Configure the CloudFront distribution behavior to forward the User-Agent header. Most Voted

upvoted 1 times

  **examaws** 6 months, 1 week ago

Selected Answer: A



Here's why A is preferred over B:

Forwarding User-Agent: Option A specifically forwards the User-Agent header from the client to the origin. This allows the backend (ALB) to determine how to respond based on the type of device making the request (mobile vs. desktop).

Custom Headers: Option B involves adding a User-Agent header as a custom header, which may not be recognized by the application unless it is explicitly programmed to handle that custom header. This could lead to the application still serving the wrong version of the site.

In summary, forwarding the User-Agent header (Option A) is a more straightforward and effective solution for ensuring that the correct version of the website is served based on the device type.

upvoted 4 times

  **ranjeetrrv** 9 months, 1 week ago

A is correct answer:

The correct action to resolve the issue of mobile users being served the desktop version of the website is:

A. Configure the CloudFront distribution behavior to forward the User-Agent header.

Explanation:

When mobile users access the website through CloudFront, the User-Agent header contains information about the device and browser being used. By configuring the CloudFront distribution behavior to forward the User-Agent header to the ALB, the ALB can use this information to serve the appropriate version of the website to mobile users

upvoted 1 times

  **Christina666** 9 months, 1 week ago

Selected Answer: A

User-Agent header

If you want CloudFront to cache different versions of your objects based on the device that a user is using to view your content, we recommend that you configure CloudFront to forward one or more of the following headers to your custom origin:

CloudFront-Is-Desktop-Viewer

CloudFront-Is-Mobile-Viewer

CloudFront-Is-SmartTV-Viewer

CloudFront-Is-Tablet-Viewer

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/RequestAndResponseBehaviorCustomOrigin.html#request-custom-user-agent-header>

upvoted 5 times

🗨️ 👤 **be9z** 9 months, 1 week ago

Option A: Configure the CloudFront distribution behavior to forward the User-Agent header.

Explanation:

When CloudFront serves content, it can use the User-Agent header from the client's request to determine the type of device (e.g., desktop or mobile). By forwarding the User-Agent header to the origin (in this case, the ALB), the origin server can respond with the appropriate version of the website based on the device type.

This ensures that mobile users receive the correct content.

upvoted 2 times

🗨️ 👤 **axdevops** 9 months, 1 week ago

Selected Answer: A

- The User-Agent header contains information about the client's device and browser. By forwarding this header, CloudFront can use it to determine whether the request is coming from a desktop or mobile device.

- Forwarding the User-Agent header allows the origin (ALB in this case) to make decisions based on the client's device type and serve the appropriate content.

"Option B is not the best choice because adding a User-Agent header to the list of origin custom headers is used for sending additional headers to the origin, not for identifying the client's device type."

Options C and D: Enabling IPv6 or updating to the dualstack endpoint does not address the mobile/desktop version problem.

upvoted 1 times

🗨️ 👤 **MohamedMarzouk** 9 months, 2 weeks ago

Answer is A

upvoted 1 times

🗨️ 👤 **64rl0** 9 months, 3 weeks ago

Selected Answer: A

Answer is A

upvoted 1 times

🗨️ 👤 **flaacko** 10 months ago

Selected Answer: A

The correct answer is A because it talks about setting the distribution behavior to make use of the User-Agent header. Option B talks about using the User-Agent header but then refers to origin settings which have nothing to do with how CloudFront processes a user's request. Option C and D are not relevant to this question.

upvoted 1 times

🗨️ 👤 **pekalyok** 1 year, 3 months ago

Selected Answer: A

A is the correct answer

upvoted 2 times

🗨️ 👤 **konieczny69** 1 year, 6 months ago

Selected Answer: A

chatgpt says A

upvoted 2 times

🗨️ 👤 **bruppp31** 1 year, 7 months ago

The correct answer is B, A is incorrect because it is FORWARDING the User-Agent. Not correct. We want to receive/use the User-Agent data to decide what content to deliver.

upvoted 4 times



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

C. Enable IPv6 on the ALB. Update the CloudFront distribution origin settings to use the dualstack endpoint.

upvoted 1 times

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

Wrong. Do you even understand what you are saying? What does IPv6 have to do here? Suggest to gain some knowledge before commenting here
upvoted 3 times

  **Benly** 1 year, 8 months ago

Selected Answer: B

B is the answer.

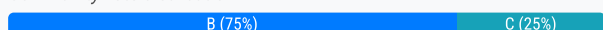
upvoted 1 times

A SysOps administrator has enabled AWS CloudTrail in an AWS account. If CloudTrail is disabled, it must be re-enabled immediately. What should the SysOps administrator do to meet these requirements WITHOUT writing custom code?

- A. Add the AWS account to AWS Organizations. Enable CloudTrail in the management account.
- B. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Apply the AWS-ConfigureCloudTrailLogging automatic remediation action.
- C. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Configure the rule to invoke an AWS Lambda function to enable CloudTrail.
- D. Create an Amazon EventBridge (Amazon CloudWatch Event) hourly rule with a schedule pattern to run an AWS Systems Manager Automation document to enable CloudTrail.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **Cyberkayu** 1 month, 2 weeks ago

Selected Answer: B

answer C required writing custom code in Lambda

upvoted 1 times

🗳️ 👤 **Xavier_Examen** 2 months, 3 weeks ago

Selected Answer: B

Option B

upvoted 1 times

🗳️ 👤 **ogum** 4 months, 3 weeks ago

Selected Answer: B

B. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Apply the AWS-ConfigureCloudTrailLogging automatic remediation action. Most Voted

upvoted 1 times

🗳️ 👤 **examaws** 6 months, 1 week ago

Selected Answer: B

Explanation:

Option B directly addresses the requirement to re-enable CloudTrail without writing custom code. By using an AWS Config rule with an automatic remediation action, it ensures that CloudTrail is enabled whenever its configuration changes, thus meeting the requirement efficiently.

Why not the others?

A: Adding the account to AWS Organizations and enabling CloudTrail in the management account does not ensure immediate re-enablement of CloudTrail in the specific account.

C: This option involves invoking a Lambda function, which implies writing custom code, contrary to the requirement.

D: While using EventBridge could work, it introduces unnecessary complexity and does not directly address the immediate need to re-enable CloudTrail without custom code.

upvoted 1 times

🗳️ 👤 **64rl0** 9 months, 3 weeks ago

Selected Answer: B

Answer is B

upvoted 1 times

🗳️ 👤 **Rabbit117** 1 year, 5 months ago

Selected Answer: B

B. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Apply the AWS-ConfigureCloudTrailLogging automatic remediation action.

upvoted 1 times

🗳️ 👤 **NAVADIYA** 1 year, 8 months ago

B. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Apply the AWS-ConfigureCloudTrailLogging automatic remediation action.

upvoted 1 times

🗲️ 👤 **arana1992** 1 year, 8 months ago

B. Create an AWS Config rule that is invoked when CloudTrail configuration changes. Apply the AWS-ConfigureCloudTrailLogging automatic remediation action.

Option B allows for automatic remediation of CloudTrail configuration changes. By creating an AWS Config rule with the AWS-ConfigureCloudTrailLogging remediation action, you can ensure that if CloudTrail is ever disabled, it will be automatically re-enabled.

Option A (adding the AWS account to AWS Organizations and enabling CloudTrail in the management account) is not directly related to re-enabling CloudTrail if it's disabled.

Option C (creating an AWS Config rule to invoke a Lambda function) would require writing custom code, which is specifically mentioned as not being allowed in the question.

Option D (creating an Amazon EventBridge rule with an Automation document) would also require custom code through AWS Systems Manager Automation documents, which is not allowed as per the question's constraints.

upvoted 1 times

🗲️ 👤 **callspace** 1 year, 9 months ago

B is correct as question clearly says WITHOUT writing custom code so C can't be correct.

upvoted 1 times

🗲️ 👤 **marcoeu** 1 year, 9 months ago

B ... <https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automatically-re-enable-aws-cloudtrail-by-using-a-custom-remediation-rule-in-aws-config.html>

upvoted 1 times

🗲️ 👤 **bakamon** 1 year, 11 months ago

Selected Answer: C

AWS-ConfigureCloudTrailLogging does not exist at all

upvoted 1 times

🗲️ 👤 **AgboolaKun** 1 year, 1 month ago

The answer is B. I did not know that AWS-ConfigureCloudTrailLogging exist in AWS Systems Manager until I checked too. You can find it in Systems Manager -> Documents, then check Automation documents box under Categories session, then you will see "AWS-ConfigureCloudTrailLogging". In fact, if you click on AWS-ConfigureCloudTrailLogging link, you will see a state machine visual that explains how to use this automation.

upvoted 1 times

🗲️ 👤 **elanelans** 1 year, 11 months ago

It does...

Login to account > AWS Systems Manager > Documents (Under shared resources)

All documents then search for key word "AWS-ConfigureCloudTrailLogging"

upvoted 5 times

🗲️ 👤 **mamila** 2 years ago

Selected Answer: C

AWS-ConfigureCloudTrailLogging does not exist, a lambda has to be called to enable CloudTrail answer is C.

upvoted 2 times

🗲️ 👤 **elanelans** 1 year, 11 months ago

It does...

Login to account > AWS Systems Manager > Documents (Under shared resources)

All documents then search for key word "AWS-ConfigureCloudTrailLogging"

upvoted 3 times

🗲️ 👤 **Gomer** 2 years, 1 month ago

I have a hard time voting for "B" just because there is no "AWS-ConfigureCloudTrailLogging" Config rule, SSM Document, SSM Runbook, SSM Automation. There is a SSM "Runbook" named "AWS-EnableCloudTrail" that I presume would make "D" work, but it seems kludgy to check hourly for something that could be automated to turn on when it's turned off with no wait period. Not sure if this is a trick question or just a poorly worded

question. "B" is wrong if you take the wording literally. If you presume they really meant to say was to use "cloudtrail-enabled" config rule, then it might be correct. But that is NOT what it says.

upvoted 2 times

🗨️ 👤 **CVDON** 2 years, 4 months ago

B But i would use SCP to prevent any disabling action. <https://aws.amazon.com/es/blogs/industries/best-practices-for-aws-organizations-service-control-policies-in-a-multi-account-environment/>

upvoted 2 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbbbbbb

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

I agree with you all that B is the answer but that remedation doesn't exist. We'll have to add it from a template and CUSTOMIZE it so ummm....

upvoted 1 times

🗨️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: B

B, Config can check and also remediate automatically.

upvoted 2 times

A company hosts its website on Amazon EC2 instances behind an Application Load Balancer. The company manages its DNS with Amazon Route 53, and wants to point its domain's zone apex to the website.

Which type of record should be used to meet these requirements?

- A. An AAAA record for the domain's zone apex
- B. An A record for the domain's zone apex
- C. A CNAME record for the domain's zone apex
- D. An alias record for the domain's zone apex

Suggested Answer: D

Reference:

<https://aws.amazon.com/route53/faqs/>

Community vote distribution

D (100%)

 **FHU** Highly Voted 3 years, 8 months ago

Letter D is correct. Route 53 supports redirection of zone apex to the ALB via alias.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html>

upvoted 12 times

 **renatom1976** Most Recent 1 month, 2 weeks ago

Selected Answer: D

D is the correct. This question is identical and I failed choosing B.


upvoted 1 times

 **ogum** 4 months, 3 weeks ago

Selected Answer: D

D. An alias record for the domain's zone apex

upvoted 1 times

 **Rabbit117** 9 months, 1 week ago

Selected Answer: D

D is the correct answer.

The DNS protocol does not allow you to create a CNAME record for the top node of a DNS namespace, also known as the zone apex. For example, if you register the DNS name example.com, the zone apex is example.com. You cannot create a CNAME record for example.com, but you can create CNAME records for www.example.com, newproduct.example.com, and so on.

Amazon Route 53 also supports alias records, which allow you to route queries to selected AWS resources, such as CloudFront distributions and Amazon S3 buckets. Aliases are similar in some ways to the CNAME record type; however, you can create an alias for the zone apex. For more information, see [Choosing between alias and non-alias records](#).

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/ResourceRecordTypes.html#CNAMEFormat>

upvoted 4 times

 **64rl0** 9 months, 3 weeks ago

Selected Answer: D

Answer is D

upvoted 1 times

 **pedroyrosa5700** 1 year, 4 months ago

D IS CORRECT

upvoted 1 times

 **bruppp31** 1 year, 7 months ago

D. The alias record is used to route the domains top node i.e zone apex i.e. example.com. You can't create a CNAME record for the top node and A and AAAA records are for routing to IPv4 and IPv6 respectively

upvoted 1 times

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

D. An alias record for the domain's zone apex

upvoted 1 times

🗨️ 👤 **nolly222** 1 year, 10 months ago

Selected Answer: D

answer is D

upvoted 1 times

🗨️ 👤 **shivangee** 1 year, 11 months ago

An A record for the domain's zone apex answer B correct answer

upvoted 1 times

🗨️ 👤 **shivangee** 1 year, 11 months ago

answer B

upvoted 1 times

🗨️ 👤 **Christina666** 1 year, 11 months ago

Selected Answer: D

The DNS protocol does not allow you to create a CNAME record for the top node of a DNS namespace, also known as the zone apex. For example, if you register the DNS name example.com, the zone apex is example.com. You cannot create a CNAME record for example.com, but you can create CNAME records for www.example.com, newproduct.example.com, and so on.

In addition, if you create a CNAME record for a subdomain, you cannot create any other records for that subdomain. For example, if you create a CNAME for www.example.com, you cannot create any other records for which the value of the Name field is www.example.com.

upvoted 1 times

🗨️ 👤 **Christina666** 1 year, 11 months ago

Unlike a CNAME record, you can create an alias record at the top node of a DNS namespace, also known as the zone apex. For example, if you register the DNS name example.com, the zone apex is example.com. You can't create a CNAME record for example.com, but you can create an alias record for example.com that routes traffic to www.example.com (as long as the record type for www.example.com is not of type CNAME).

upvoted 3 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

dddddddd

upvoted 1 times

🗨️ 👤 **mlantonis2** 2 years, 7 months ago

Ans: D

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: D

upvoted 1 times

🗨️ 👤 **Anthony053** 2 years, 8 months ago

Alias: Used to map DNS record to Amazon S3 bucket static website.

upvoted 1 times

🗨️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: D

D. Cannot create CNAME record for zone apex. AAAA and A are address records, which translates website names to IPv6 and IPv4 addresses respectively.

upvoted 2 times

A company must ensure that any objects uploaded to an S3 bucket are encrypted.

Which of the following actions will meet this requirement? (Choose two.)

- A. Implement AWS Shield to protect against unencrypted objects stored in S3 buckets.
- B. Implement Object access control list (ACL) to deny unencrypted objects from being uploaded to the S3 bucket.
- C. Implement Amazon S3 default encryption to make sure that any object being uploaded is encrypted before it is stored.
- D. Implement Amazon Inspector to inspect objects uploaded to the S3 bucket to make sure that they are encrypted.
- E. Implement S3 bucket policies to deny unencrypted objects from being uploaded to the buckets.

Suggested Answer: CE

Reference:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/acl-overview.html#sample-acl>

Access control list (ACL) overview

[PDF](#) | [RSS](#)

Amazon S3 access control lists (ACLs) enable you to manage access to buckets and objects. Each bucket and object has an ACL attached to it as a subresource. It defines which AWS accounts or groups are granted access and the type of access. When a request is received against a resource, Amazon S3 checks the corresponding ACL to verify that the requester has the necessary access permissions.

By default, when another AWS account uploads an object to your S3 bucket, that account (the object writer) owns the object, has access to it, and can grant other users access to it through ACLs. You can use Object Ownership to change this default behavior so that ACLs are disabled and you, as the bucket owner, automatically own every object in your bucket. As a result, access control for your data is based on policies, such as IAM policies, S3 bucket policies, virtual private cloud (VPC) endpoint policies, and AWS Organizations service control policies (SCPs).

Community vote distribution

CE (100%)

 **CVDON** Highly Voted 2 years, 4 months ago

C and E.

S3 encryption and

S3 bucket policy with deny S3 Put request without x-amz-server-side-encryption header

upvoted 9 times

 **mimahmed_awseducate** 1 year, 7 months ago

Right Answer

upvoted 1 times

 **ogum** Most Recent 4 months, 3 weeks ago

Selected Answer: CE

Answer is CE

upvoted 1 times

 **64r10** 9 months, 3 weeks ago

Selected Answer: CE

Answer is CE

upvoted 1 times

 **NAVADIYA** 1 year, 8 months ago

C. Implement Amazon S3 default encryption to make sure that any object being uploaded is encrypted before it is stored. Most Voted

E. Implement S3 bucket policies to deny unencrypted objects from being uploaded to the buckets.

upvoted 2 times

🗳️ 👤 **CVDON** 2 years, 4 months ago

<https://aws.amazon.com/es/blogs/aws/amazon-s3-encrypts-new-objects-by-default/>

upvoted 1 times

🗳️ 👤 **BietTuot** 2 years, 6 months ago

Selected Answer: CE

C and E

upvoted 2 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

ccc eee

upvoted 2 times

🗳️ 👤 **mlantonis2** 2 years, 7 months ago

Selected Answer: CE

Ans: CE

upvoted 2 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: CE

upvoted 1 times

🗳️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: CE

C and E

upvoted 1 times

🗳️ 👤 **nakikoo** 2 years, 9 months ago

CE correct, default encryption is a feature you can enable and disable in S3, it encrypt the data when entered S3 and decrypt whenever people retrieve data...server-side encryption is data encrypted as it is before entering an S3..

upvoted 1 times

🗳️ 👤 **MikeyJ** 3 years, 1 month ago

Poorly worded question as encrypting objects before uploading would use client side encryption. C&E seem the most likely answers, as ACLs can't prevent the uploading of unencrypted objects.

upvoted 2 times

🗳️ 👤 **by116549** 3 years, 1 month ago

Sorry @Finger41 and @Mecdrox I am bit confused by C as the question states:

"verify that all items uploaded to an S3 bucket are encrypted prior to uploading them"

Option C from what I can see states:

"With Amazon S3 default encryption, you can set the default encryption behavior for an S3 bucket so that all new objects are encrypted when they are stored in the bucket. The objects are encrypted using server-side encryption with either Amazon S3-managed keys (SSE-S3) or AWS KMS keys stored in AWS Key Management Service (AWS KMS) (SSE-KMS)."

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html>

Does the data not need to be encrypted prior to being uploaded?

upvoted 1 times

🗳️ 👤 **Finger41** 3 years, 1 month ago

Its encrypted at the time of writing to disk. :). Ensures all objects are encrypted when data is stored in S3, if using Amazon S3 default encryption ie server side encryption.

Looking at an extension of your link : <https://docs.aws.amazon.com/AmazonS3/latest/userguide/serv-side-encryption.html>

Amazon S3 encrypts your data at the object level as it writes it to disks in its data centers and decrypts it for you when you access it.

upvoted 2 times

🗳️ 👤 **Finger41** 3 years, 1 month ago

Selected Answer: CE


C & E - <https://aws.amazon.com/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/>

upvoted 3 times

  **MikeyJ** 3 years ago

"In order to enforce object encryption, create an S3 bucket policy that denies any S3 Put request that does not include the x-amz-server-side-encryption header. There are two possible values for the x-amz-server-side-encryption header: AES256, which tells S3 to use S3-managed keys, and aws:kms, which tells S3 to use AWS KMS-managed keys."

upvoted 1 times

  **Mecdrex** 3 years, 2 months ago

Selected Answer: CE

C and E are correct

upvoted 1 times

A company has a stateful web application that is hosted on Amazon EC2 instances in an Auto Scaling group. The instances run behind an Application Load Balancer (ALB) that has a single target group. The ALB is configured as the origin in an Amazon CloudFront distribution. Users are reporting random logouts from the web application.

Which combination of actions should a SysOps administrator take to resolve this problem? (Choose two.)

- A. Change to the least outstanding requests algorithm on the ALB target group.
- B. Configure cookie forwarding in the CloudFront distribution cache behavior.
- C. Configure header forwarding in the CloudFront distribution cache behavior.
- D. Enable group-level stickiness on the ALB listener rule.
- E. Enable sticky sessions on the ALB target group.

Suggested Answer: CE

Community vote distribution

BE (100%)

 **jkwek** Highly Voted 9 months, 1 week ago

Answer is B and E.

Refer url :

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

You can configure each cache behavior to do one of the following:

Forward all cookies to your origin – CloudFront includes all cookies sent by the viewer when it forwards requests to the origin.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/sticky-sessions.html>

By default, an Application Load Balancer routes each request independently to a registered target based on the chosen load-balancing algorithm.

upvoted 27 times

 **Huy** Highly Voted 3 years, 7 months ago

B & E is correct answer. Enable sticky session + forward Cookie because ALB sticky session works based on cookie.

upvoted 13 times

 **mimahmed_awseducate** 1 year, 7 months ago

I also agree. I the correct answer should be B and E.

upvoted 1 times

 **ogum** Most Recent 4 months, 3 weeks ago

Selected Answer: BE

Answer is B and E, guys.

upvoted 1 times

 **ft_cloud** 9 months, 1 week ago

Selected Answer: BE

Answer is B and E, guys.

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


upvoted 1 times

 **64rl0** 9 months, 3 weeks ago

Selected Answer: BE

Answer is BE

upvoted 2 times

 **pekalyok** 1 year, 3 months ago

Selected Answer: BE

b and e

upvoted 1 times

🗳️ 👤 **NAVADIYA** 1 year, 8 months ago

C. Configure header forwarding in the CloudFront distribution cache behavior.

E. Enable sticky sessions on the ALB target group.

upvoted 1 times

🗳️ 👤 **shivangee** 1 year, 11 months ago

Answer B and E -- Configure cookie forwarding in the CloudFront distribution cache behavior.

upvoted 1 times

🗳️ 👤 **Christina666** 1 year, 11 months ago

Selected Answer: BE

D. Enabling group-level stickiness on the ALB listener rule does not exist as an option. While ALB supports sticky sessions through target group settings, there is no concept of group-level stickiness at the listener rule level.

upvoted 3 times

🗳️ 👤 **Andrew_A** 2 years ago

Option C (header forwarding in the CloudFront distribution) could potentially help if the application uses headers to maintain session state. However, the question suggests that this is a cookie-based issue. Hence Options B & E are correct.

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: BE

bbbbbbb and eeeeeee

upvoted 1 times

🗳️ 👤 **mlantonis2** 2 years, 7 months ago

Ans: BE

upvoted 2 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: BE

upvoted 1 times

🗳️ 👤 **nakikoo** 2 years, 9 months ago

im sticking with CE with this one...again might be wrong, but logically, the questions says "Stateful" web application, meaning it comes from one host, header can be set from the one stateful web application to route to only one destination...configuring cookie based might overload the process, which explains the logout issue...but i maybe wrong, just a suggestion

Header based caching:

<https://aws.amazon.com/premiumsupport/knowledge-center/configure-cloudfront-to-forward-headers/>

cookie based caching:

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

upvoted 1 times

🗳️ 👤 **mohamedba** 2 years, 11 months ago

Cookie forwarding for CloudFront and Stick Session for the ALB

upvoted 1 times

🗳️ 👤 **Mikilo** 3 years, 1 month ago

Selected Answer: BE

B and E is correct answer

upvoted 1 times

🗳️ 👤 **Finger41** 3 years, 1 month ago

Selected Answer: BE

B and E : <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Cookies.html> +

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/sticky-sessions.html>

upvoted 3 times

A company is running a serverless application on AWS Lambda. The application stores data in an Amazon RDS for MySQL DB instance. Usage has steadily increased, and recently there have been numerous "too many connections" errors when the Lambda function attempts to connect to the database. The company already has configured the database to use the maximum max_connections value that is possible.


What should a SysOps administrator do to resolve these errors?

- A. Create a read replica of the database. Use Amazon Route 53 to create a weighted DNS record that contains both databases.
- B. Use Amazon RDS Proxy to create a proxy. Update the connection string in the Lambda function.
- C. Increase the value in the max_connect_errors parameter in the parameter group that the database uses.
- D. Update the Lambda function's reserved concurrency to a higher value.

Suggested Answer: B

Community vote distribution

B (100%)

 **binhdt2611** Highly Voted 3 years, 8 months ago

Answer is B

Check "Database proxy for Amazon RDS" section in the link to see how RDS proxy help Lambda handle huge connections to RDS MySQL

<https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>
upvoted 18 times

 **jkwek** Highly Voted 9 months, 1 week ago

Answer is B.

<https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>

RDS Proxy acts as an intermediary between your application and an RDS database. RDS Proxy establishes and manages the necessary connection pools to your database so that your application creates fewer database connections.


Your Lambda functions interact with RDS Proxy instead of your database instance. It handles the connection pooling necessary for scaling many simultaneous connections created by concurrent Lambda functions. This allows your Lambda applications to reuse existing connections, rather than creating new connections for every function invocation.

upvoted 15 times

 **ogum** Most Recent 4 months, 3 weeks ago

Selected Answer: B

B. Use Amazon RDS Proxy to create a proxy. Update the connection string in the Lambda function. Most Voted
upvoted 1 times

 **Dinya_jui** 5 months, 4 weeks ago

Selected Answer: B

RDS proxy for connections management with AWS Lambda.

<https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>
upvoted 1 times


 **james2033** 1 year, 3 months ago

Selected Answer: B

RDS proxy for connections management with AWS Lambda.
upvoted 1 times

 **NAVADIYA** 1 year, 8 months ago

B. Use Amazon RDS Proxy to create a proxy. Update the connection string in the Lambda function.
upvoted 1 times

 **CVDON** 2 years, 4 months ago



B. Improve efficiency in connection management and reuse
upvoted 1 times

 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbbbbbb

upvoted 1 times



  **Finger41** 3 years, 1 month ago

Selected Answer: B

B - <https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>

Solution straight out of the docs....

upvoted 4 times

  **pglag** 3 years, 3 months ago

Selected Answer: B

its B for sure



upvoted 2 times

  **mmmmm12451** 3 years, 3 months ago

Selected Answer: B



B for sure

upvoted 2 times

  **VTHOR** 3 years, 3 months ago

B, for sure


upvoted 1 times

  **roka_ua** 3 years, 3 months ago

Selected Answer: B

Should be B

upvoted 1 times

  **MrkJobs** 3 years, 5 months ago



Selected Answer: B

B

"You can use RDS Proxy for any application that makes SQL calls to your database. But in the context of serverless, we focus on how this improves the Lambda experience. The proxy handles all database traffic that normally flows from your Lambda functions directly to the database."

<https://aws.amazon.com/blogs/compute/using-amazon-rds-proxy-with-aws-lambda/>

upvoted 3 times

  **doc_nta** 3 years, 6 months ago

Answer is B



upvoted 1 times

  **ngthien041292** 3 years, 7 months ago

Selected Answer: B

Vote B

upvoted 2 times

  **lei00** 3 years, 7 months ago

I think the answer should be B.

D -- the concurrency is more towards the services invoke Lambda, instead of the question here "too many connections to RDS"

<https://aws.amazon.com/tw/blogs/compute/managing-aws-lambda-function-concurrency/>

upvoted 2 times

A SysOps administrator is deploying an application on 10 Amazon EC2 instances. The application must be highly available. The instances must be placed on distinct underlying hardware.

What should the SysOps administrator do to meet these requirements?

- A. Launch the instances into a cluster placement group in a single AWS Region.
- B. Launch the instances into a partition placement group in multiple AWS Regions.
- C. Launch the instances into a spread placement group in multiple AWS Regions.
- D. Launch the instances into a spread placement group in a single AWS Region.

Suggested Answer: B

Community vote distribution

D (100%)

🗳️ 👤 **Azaad78** Highly Voted 👍 3 years, 9 months ago

D - I think. Keep in mind placement groups can't cross regions.

upvoted 26 times

🗳️ 👤 **jipark** 1 year, 10 months ago

support multi-AZ, not multi-Region

upvoted 2 times

🗳️ 👤 **MrkJobs** Highly Voted 👍 9 months, 1 week ago

Selected Answer: D

answer D:

Spread works only in 1 region.

Max instance per availability zone 7. if 3 availability zone = 21 Instance. (Question=10 instance)

"A spread placement group is a group of instances that are each placed on distinct racks, with each rack having its own network and power source."

All instance are placed on different racks for a maximum of 7 instance for Availability Zone.

"A spread placement group can span multiple Availability Zones in the same Region. You can have a maximum of seven running instances per Availability Zone per group."

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

upvoted 6 times

🗳️ 👤 **64rl0** Most Recent 🕒 9 months, 3 weeks ago

Selected Answer: D

Answer is D

upvoted 1 times

🗳️ 👤 **pekalyok** 1 year, 3 months ago

Selected Answer: D

D is the answer

upvoted 2 times

🗳️ 👤 **konieczny69** 1 year, 6 months ago

Selected Answer: D

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

A spread placement group is a group of instances that are each placed on distinct hardware.

upvoted 1 times

🗳️ 👤 **NAVADIYA** 1 year, 8 months ago

B. Launch the instances into a partition placement group in multiple AWS Regions.

upvoted 1 times

🗳️ 👤 **arana1992** 1 year, 8 months ago

D. Launch the instances into a spread placement group in a single AWS Region.

A spread placement group is a logical grouping of instances that are placed on distinct underlying hardware to reduce the risk of simultaneous

failures due to hardware issues. This ensures high availability.

In this scenario, the requirement is to have the instances on distinct underlying hardware, and since the application needs to be highly available, you should launch them in a single AWS Region to keep them close for low latency.

Options A and B are incorrect because they involve multiple AWS Regions, which is not necessary for the given requirements. Option C is incorrect because it suggests using a spread placement group in multiple AWS Regions, which is not supported.

upvoted 3 times

🗳️ 👤 **habros** 1 year, 10 months ago

Selected Answer: D

Placement groups are within AZs, hence one AZ correlates to only one Region

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

Spread placement for unique hardware across the same AZ for ultimate redundancy

Cluster placement uses hardware within the same AZ for faster networking

Partition placement deploys equal sets of hardware across unique racks for redundancy

upvoted 4 times

🗳️ 👤 **shivangee** 1 year, 11 months ago

D. Launch the instances into a spread placement group in a single AWS Region is right answer.

upvoted 1 times

🗳️ 👤 **nizammusasoac02** 1 year, 11 months ago

Selected Answer: D

What is the difference between spread placement group and partition group?

The spread placement group setup may be similar to partition placement groups, but the main difference is that partition placement groups are made of several instances on each partition, while spread groups are just single individual instances spread through different racks or AZs.

upvoted 2 times

🗳️ 👤 **Bhrino** 2 years, 3 months ago

Selected Answer: D

We know its spread placement because this is associated with high availability. But its D because they cannot travel across regions

upvoted 2 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

dddddd

upvoted 1 times

🗳️ 👤 **pravinb** 2 years, 7 months ago

B incorrect - you cant launch them in multiple AWS regions.. Only D

upvoted 2 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: D

upvoted 1 times

🗳️ 👤 **yeanningmedal71** 2 years, 11 months ago

Selected Answer: D

Spread – strictly places a small group of instances across distinct underlying hardware to reduce correlated failures. A rack spread placement group can span multiple Availability Zones in the same Region. For rack spread level placement groups, you can have a maximum of seven running instances per Availability Zone per group.

upvoted 1 times

🗳️ 👤 **mohamedba** 2 years, 11 months ago

D.

Spread in single region is the answer (max 7 instances per AZ). Not that placement groups can't cross regions.

upvoted 1 times

🗳️ 👤 **ceeee** 3 years ago

Selected Answer: D

D definitely, because we need ec2s in separated racks and because placement groups do not cross regions

upvoted 3 times

A SysOps administrator is troubleshooting an AWS CloudFormation template whereby multiple Amazon EC2 instances are being created. The template is working in us-east-1, but it is failing in us-west-2 with the error code:

AMI [ami-12345678] does not exist

How should the Administrator ensure that the AWS CloudFormation template is working in every region?

- A. Copy the source region's Amazon Machine Image (AMI) to the destination region and assign it the same ID.
- B. Edit the AWS CloudFormation template to specify the region code as part of the fully qualified AMI ID.
- C. Edit the AWS CloudFormation template to offer a drop-down list of all AMIs to the user by using the AWS::EC2::AMI::ImageID control.
- D. Modify the AWS CloudFormation template by including the AMI IDs in the `Mappings` section. Refer to the proper mapping within the template for the proper AMI ID.

Suggested Answer: D

Community vote distribution

D (100%)

 **Goozian** Highly Voted 9 months, 1 week ago

Selected Answer: D

Parameters:

EnvironmentType:

Description: The environment type

Type: String

Default: test

AllowedValues:

- prod

- test

ConstraintDescription: must be a prod or test

Mappings:

RegionAndInstanceTypeToAMIID:

us-east-1:

test: "ami-8ff710e2"

prod: "ami-f5f41398"

us-west-2:

test: "ami-eff1028f"

prod: "ami-d0f506b0"

...other regions and AMI IDs...

Resources:

...other resources...


Outputs:

TestOutput:

Description: Return the name of the AMI ID that matches the region and environment type keys

Value: !FindInMap [RegionAndInstanceTypeToAMIID, !Ref "AWS::Region", !Ref EnvironmentType]

upvoted 11 times

 **Finger41** Highly Voted 3 years, 1 month ago

Selected Answer: D

D - <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html>

upvoted 6 times

 **64rl0** Most Recent 9 months, 3 weeks ago

Selected Answer: D



Answer is D

upvoted 1 times

  **NAVADIYA** 1 year, 8 months ago

D. Modify the AWS CloudFormation template by including the AMI IDs in the `Mappings` section. Refer to the proper mapping within the template for the proper AMI ID.

upvoted 1 times

  **nolly222** 1 year, 10 months ago

Selected Answer: D

ddddddddd



upvoted 1 times

  **gasgasjb** 2 years, 6 months ago

Selected Answer: D

D. Mappings Section.

upvoted 2 times

  **jipark** 1 year, 10 months ago

"Refer To" is key word

upvoted 1 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddddd

upvoted 1 times

  **Liongeek** 2 years, 7 months ago

Ans: D


upvoted 1 times

  **Surferbolt** 2 years, 8 months ago

Selected Answer: D

D. I may be wrong in the details, but if I recall correctly, AMI IDs are different for different Region, and Mappings in CloudFormation template can solve the problem described.

upvoted 3 times

  **Mecdrex** 3 years, 2 months ago

Selected Answer: D

I vote D. Mappings is the solution.

upvoted 3 times

A SysOps administrator is provisioning an Amazon Elastic File System (Amazon EFS) file system to provide shared storage across multiple Amazon EC2 instances. The instances all exist in the same VPC across multiple Availability Zones. There are two instances in each Availability Zone. The SysOps administrator must make the file system accessible to each instance with the lowest possible latency.

Which solution will meet these requirements?

- A. Create a mount target for the EFS file system in the VPC. Use the mount target to mount the file system on each of the instances.
- B. Create a mount target for the EFS file system in one Availability Zone of the VPC. Use the mount target to mount the file system on the instances in that Availability Zone. Share the directory with the other instances.
- C. Create a mount target for each instance. Use each mount target to mount the EFS file system on each respective instance.
- D. Create a mount target in each Availability Zone of the VPC. Use the mount target to mount the EFS file system on the instances in the respective Availability Zone.

Suggested Answer: D

Reference:

<https://docs.aws.amazon.com/efs/latest/ug/accessing-fs.html>

Creating and managing mount targets

[PDF](#) | [RSS](#)

After you create an Amazon EFS file system, you can create mount targets. For Amazon EFS file systems that use Standard storage classes, you can create a mount target in each Availability Zone in an AWS Region. For EFS file systems that use One Zone storage classes, you can only create a single mount target in the same Availability Zone as the file system. Then you can mount the file system on compute instances, including Amazon EC2, Amazon ECS, and AWS Lambda in your virtual private cloud (VPC).

The following diagram shows an Amazon EFS file system that uses Standard storage classes, with mount targets created in all Availability Zones in the VPC.

Community vote distribution

D (100%)

🗳️ 👤 **Pisces225** 1 month, 3 weeks ago

Selected Answer: D

The reason it's not C: "You can create only one mount target per Availability Zone."

<https://docs.aws.amazon.com/efs/latest/ug/accessing-fs.html>

upvoted 2 times

🗳️ 👤 **Dinya_jui** 2 months, 1 week ago

Selected Answer: D

mount target to AZ and then AZ to respective EC2 instance.

upvoted 1 times

🗳️ 👤 **jepwi** 2 months, 2 weeks ago

Selected Answer: D

For High Availability

upvoted 1 times

🗳️ 👤 **64rl0** 9 months, 3 weeks ago

Selected Answer: D

Answer is D

upvoted 2 times

🗳️ 👤 **Mangesh_XI_mumbai** 1 year, 7 months ago

Selected Answer: D

mount target to AZ and then AZ to respective EC2 instance.

upvoted 2 times

🗳️ 👤 **jipark** 1 year, 10 months ago

Selected Answer: D

mount target per 'AZ' - not VPC, nor Instance

upvoted 3 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddd

upvoted 1 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: D

upvoted 1 times

🗳️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: D

D. You can create an EFS mount target down to AZ level only.

upvoted 2 times

🗳️ 👤 **221898** 3 years ago

Selected Answer: D

Correct Answer: D ☐

upvoted 1 times

🗳️ 👤 **Finger41** 3 years, 1 month ago

Selected Answer: D

d - <https://docs.aws.amazon.com/efs/latest/ug/accessing-fs.html>

upvoted 3 times

🗳️ 👤 **vjt** 3 years, 1 month ago

It is D.

A mount target provides an IP address for an NFSv4 endpoint at which you can mount an Amazon EFS file system. You mount your file system using its Domain Name Service (DNS) name, which resolves to the IP address of the EFS mount target in the same Availability Zone as your EC2 instance. You can create one mount target in each Availability Zone in an AWS Region. If there are multiple subnets in an Availability Zone in your VPC, you create a mount target in one of the subnets. Then all EC2 instances in that Availability Zone share that mount target.

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

upvoted 4 times

🗳️ 👤 **Mecdrex** 3 years, 2 months ago

Selected Answer: D

I vote D for multi AZ requirement.

upvoted 1 times

A SysOps administrator has successfully deployed a VPC with an AWS CloudFormation template. The SysOps administrator wants to deploy the same template across multiple accounts that are managed through AWS Organizations. Which solution will meet this requirement with the LEAST operational overhead?

- A. Assume the OrganizationAccountAccessRole IAM role from the management account. Deploy the template in each of the accounts.
- B. Create an AWS Lambda function to assume a role in each account. Deploy the template by using the AWS CloudFormation CreateStack API call.
- C. Create an AWS Lambda function to query for a list of accounts. Deploy the template by using the AWS CloudFormation CreateStack API call.
- D. Use AWS CloudFormation StackSets from the management account to deploy the template in each of the accounts.

Suggested Answer: D

Reference:

<https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws-organization/>

Community vote distribution

D (100%)

 **Hello23** Highly Voted 2 years, 10 months ago

Selected Answer: D

AWS CloudFormation StackSets extends the capability of stacks by enabling you to create, update, or delete stacks across multiple accounts and AWS Regions
upvoted 8 times

 **jepwi** Most Recent 2 months, 2 weeks ago

Selected Answer: D

Ans: D
upvoted 1 times

 **Kkosha** 6 months, 1 week ago

Selected Answer: D

Want to use same template in Multiple Accounts.
upvoted 1 times

 **64rl0** 9 months, 3 weeks ago

Selected Answer: D

Answer is D
upvoted 1 times

 **airraid2010** 1 year, 6 months ago


Answer: D
<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>

With service-managed permissions, you can deploy stack instances to accounts managed by AWS Organizations. Using this permissions model, you don't have to create the necessary IAM roles; StackSets creates the IAM roles on your behalf. With this model, you can also turn on automatic deployments to accounts that you add to your organization in the future.
upvoted 3 times

 **Mangesh_XI_mumbai** 1 year, 7 months ago

Selected Answer: D



d is correct, stacksets using cloudformation through org
upvoted 1 times

 **alexiscloud** 1 year, 8 months ago

StackSet consists of a template, parameters and the list of accounts/regions where stacks will be deployed.

D

upvoted 1 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: D

dddddd

upvoted 2 times

  **mlantonis2** 2 years, 7 months ago


D StackSets

upvoted 1 times

  **Liongeek** 2 years, 7 months ago

Ans: D

upvoted 1 times

  **Surferbolt** 2 years, 8 months ago

D. CloudFormation StackSets is designed for scenarios like that.

upvoted 4 times

  **andrerkn** 2 years, 9 months ago

Selected Answer: D

D is the correct

upvoted 1 times

A company is running distributed computing software to manage a fleet of 20 Amazon EC2 instances for calculations. The fleet includes 2 control nodes and 18 task nodes to run the calculations. Control nodes can automatically start the task nodes. Currently, all the nodes run on demand. The control nodes must be available 24 hours a day, 7 days a week. The task nodes run for 4 hours each day. A SysOps administrator needs to optimize the cost of this solution. Which combination of actions will meet these requirements? (Choose two.)

- A. Purchase EC2 Instance Savings Plans for the control nodes.
- B. Use Dedicated Hosts for the control nodes.
- C. Use Reserved Instances for the task nodes.
- D. Use Spot Instances for the control nodes. Use On-Demand Instances if there is no Spot availability.
- E. Use Spot Instances for the task nodes. Use On-Demand Instances if there is no Spot availability.

Suggested Answer: CE

Community vote distribution

AE (77%)

CE (23%)

  **CVDON** Highly Voted 2 years, 4 months ago

A and E, but i would choose RI for control nodes as a better option.
upvoted 12 times

  **Gomer** Highly Voted 2 years, 3 months ago

Can't be "C" and "E" because they are contradictory. Both of these answers suggest two different solutions for "task nodes". People who are choosing "C" are thinking its directed to "control nodes", but it is not. Answer has to be A and E. I'd agree that control nodes should be reserved instances, (or "A") but that is not what "C" is proposing. It's a trick question.
upvoted 6 times

  **piavik** 2 years, 2 months ago

C is not correct
C. Use Reserved Instances for the TASK nodes.
not CONTROL nodes
upvoted 5 times

  **jepwi** Most Recent 2 months, 2 weeks ago

Selected Answer: AE
Key word is Optimize
upvoted 1 times

  **numark** 6 months, 3 weeks ago

Selected Answer: AE
Why would you purchase reserved for the task nodes that only need to run for 4 hours? I hope you are not on the AWS team in my company.
upvoted 1 times

  **64rl0** 9 months, 3 weeks ago

Selected Answer: AE
Answer is AE
upvoted 1 times

  **stallionaws** 10 months ago

A and E are perfect for this
Savings Plans for Control Nodes (Option A):

Cost-Effectiveness: Savings Plans provide substantial cost savings for workloads that need to run continuously. Since control nodes need to be available 24/7, Savings Plans are an effective way to reduce costs while maintaining consistent availability.

Flexibility: Savings Plans are more flexible than Reserved Instances as they apply to any instance usage within a specific region.

Spot Instances for Task Nodes (Option E):

Cost Reduction: Spot Instances can offer up to 90% savings over On-Demand prices, making them an excellent choice for non-critical, short-term workloads like the task nodes.

Fallback Mechanism: By falling back to On-Demand Instances when Spot Instances are unavailable, you can ensure that the task nodes still execute as required without significant cost implications.

upvoted 1 times

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

Selected Answer: AE

how C&E, A&E, what about control nodes.

upvoted 1 times

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

A&E seems correct.

upvoted 1 times

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

These are not 'good' options in my opinion.

Reserved instances for control nodes and spot block for task nodes would be more ideal.

upvoted 2 times

🗨️ 👤 **callspace** 1 year, 9 months ago

Selected Answer: AE

My Apologies. Overlooked Task nodes are suppose to be spot instances and control nodes needs to be 24/7. AE It is.

upvoted 1 times

🗨️ 👤 **callspace** 1 year, 9 months ago

Selected Answer: CE

As we know from the question - The control nodes must be available 24 hours a day, 7 days a week. CE appears to be tight choice.

Reserved Instances can provide higher savings than Savings Plan for predictable workloads with consistent usage patterns.

Savings Plan can provide higher savings than Reserved Instances for workloads with varying usage patterns or instance types.

<https://medium.com/@ekantmate/savings-plan-vs-reserved-instances-which-one-to-select-cdb8dba9facd>

upvoted 3 times

🗨️ 👤 **flaacko** 9 months, 4 weeks ago

CE is not the most cost effective choice because option C recommends using reserved instances for the task nodes which run only 4 hours a day. By using reserved instances you are committing to paying for a on-demand instances on a per hour basis for a term of 1 or 3 years. Imagine paying for something every 24 hours in a day for a year or 3 years when you only need to use it for 4 hours each. Doesn't sound cost effective right? It will make sense if you were using reserved instances for the control nodes that need to run 24/7. For the control nodes however, saving plans are still a better choice as they offer more flexibility over reserved instances.

upvoted 1 times

🗨️ 👤 **sisover** 1 year, 11 months ago

Selected Answer: AE

I disagree that C is the correct one.

How can we ask them to correct this one ?

upvoted 1 times

🗨️ 👤 **shivangee** 1 year, 11 months ago

A and E A. Purchase EC2 Instance Savings Plans for the control nodes.

upvoted 2 times

🗨️ 👤 **jipark** 1 year, 10 months ago

thanks a lot.

now I got to know "control node" and "task nodes" - reserved

upvoted 1 times

🗨️ 👤 **nizammasoac02** 1 year, 11 months ago

Selected Answer: AE

A and E

upvoted 2 times

🗨️ 👤 **BietTuot** 2 years, 6 months ago

Selected Answer: AE

A and E



upvoted 1 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: AE

A and E

upvoted 1 times

  **pravinb** 2 years, 7 months ago

A and E make more sense if we compare all other options.. we cant use spot instances for control nodes..

upvoted 1 times

  **mlantonis2** 2 years, 7 months ago

A and E.

upvoted 1 times

A company is supposed to receive a data file every hour in an Amazon S3 bucket. An S3 event notification invokes an AWS Lambda function each time a file arrives. The function processes the data for use by an application.

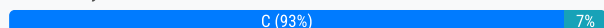
The application team notices that sometimes the file does not arrive. The application team wants to receive a notification whenever the file does not arrive.

What is the MOST operationally efficient solution that meets these requirements?

- A. Add an S3 Lifecycle rule on the S3 bucket with a scope that is limited to objects that were created in the last hour. Configure another S3 event notification to be invoked by the lifecycle transition when the number of objects transitioned is zero. Publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to notify the application team.
- B. Configure another S3 event notification to invoke a Lambda function that posts a message to an Amazon Simple Queue Service (Amazon SQS) queue. Create an Amazon CloudWatch alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to notify the application team when the ApproximateAgeOfOldestMessage metric of the queue is greater than 1 hour.
- C. Create an Amazon CloudWatch alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to alert the application team when the Invocations metric of the Lambda function is zero for an hour. Configure the alarm to treat missing data as breaching.
- D. Create a new Lambda function to get the timestamp of the newest file in the S3 bucket. If the timestamp is more than 1 hour ago, publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to notify the application team. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the new function hourly.

Suggested Answer: C

Community vote distribution



haxaffee Highly Voted 2 years, 10 months ago

Selected Answer: C

C sounds like the best option. Tested it as well.

upvoted 9 times

Nikhil3098765 Most Recent 1 month, 2 weeks ago

Selected Answer: C

Answer is C

upvoted 1 times

64rl0 9 months, 3 weeks ago

Selected Answer: C

Answer is C

upvoted 2 times

tamng 1 year, 6 months ago

C. Create an Amazon CloudWatch alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic to alert the application team when the Invocations metric of the Lambda function is zero for an hour. Configure the alarm to treat missing data as breaching.

upvoted 1 times

habros 1 year, 10 months ago

Selected Answer: C

No point sending to other Lambda functions when it cant even trigger the actual Lambda function! C is the most sane answer.

upvoted 2 times

thetnyeinmoe 2 years, 1 month ago

Selected Answer: D

This solution is the most operationally efficient as it uses a simple and straightforward approach to check if the data file is not arriving on time. The new Lambda function can be created with minimal configuration, and the Amazon CloudWatch Events rule can be set up to trigger the function hourly. If the timestamp of the newest file in the S3 bucket is more than an hour ago, the function will publish a message to an Amazon SNS topic to notify the application team. This approach avoids complex configurations and dependencies on other AWS services, making it easier to manage and maintain over time.

upvoted 1 times

🗨️ 👤 **ad45** 1 year, 9 months ago

You chose the D answer, we're talking about "most operationally" solution and the answer D uses 3 services instead of 2 like the answer C. In my opinion, the good answer is the C one.

upvoted 8 times

🗨️ 👤 **CVDON** 2 years, 4 months ago

C simple and clean as f*

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: C

cccccc

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: C

upvoted 1 times

🗨️ 👤 **Atown** 2 years, 7 months ago

Selected Answer: C

I vote C

upvoted 1 times

A company recently acquired another corporation and all of that corporation's AWS accounts. A financial analyst needs the cost data from these accounts. A

SysOps administrator uses Cost Explorer to generate cost and usage reports. The SysOps administrator notices that "No Tagkey" represents 20% of the monthly cost.

What should the SysOps administrator do to tag the "No Tagkey" resources?

- A. Add the accounts to AWS Organizations. Use a service control policy (SCP) to tag all the untagged resources.
- B. Use an AWS Config rule to find the untagged resources. Set the remediation action to terminate the resources.
- C. Use Cost Explorer to find and tag all the untagged resources.
- D. Use Tag Editor to find and tag all the untagged resources.

Suggested Answer: D

Community vote distribution

D (100%)

 **Surferbolt** Highly Voted 2 years, 8 months ago

Selected Answer: D

D.

"You can add tags to resources when you create the resource. You can use the resource's service console or API to add, change, or remove those tags one resource at a time. To add tags to—or edit or delete tags of—multiple resources at once, use Tag Editor. With Tag Editor, you search for the resources that you want to tag, and then manage tags for the resources in your search results."

<https://docs.aws.amazon.com/ARG/latest/userguide/tag-editor.html>

upvoted 7 times

 **jepwi** Most Recent 2 months, 2 weeks ago

Selected Answer: D

Tag Editor is for mass editing tag Keys and Values of all the resources in an Account

upvoted 1 times

 **64rl0** 9 months, 3 weeks ago

Selected Answer: D

Answer is D

upvoted 1 times

 **habros** 1 year, 10 months ago

Selected Answer: D

D for this.

Would have chosen B if it offers tag remediation. But in this case, it is a trick as it terminates any resources. This would be detrimental.

upvoted 3 times

 **Liongeek** 2 years, 7 months ago


Ans: D

upvoted 1 times

 **Onimole** 2 years, 9 months ago

Correct!

upvoted 1 times

 **haxaffee** 2 years, 10 months ago

Selected Answer: D

Tag Editor as part of AWS Resource Groups works for this. D.

upvoted 1 times

While setting up an AWS managed VPN connection, a SysOps administrator creates a customer gateway resource in AWS. The customer gateway device resides in a data center with a NAT gateway in front of it.

What address should be used to create the customer gateway resource?

- A. The private IP address of the customer gateway device
- B. The MAC address of the NAT device in front of the customer gateway device
- C. The public IP address of the customer gateway device
- D. The public IP address of the NAT device in front of the customer gateway device

Suggested Answer: D

Reference:

<https://docs.aws.amazon.com/vpn/latest/s2svpn/cgw-options.html>

Community vote distribution

D (100%)

🗳️ 👤 **moonwalkeryj** Highly Voted 2 years, 4 months ago

D

If your customer gateway device is behind a network address translation (NAT) device, use the IP address of your NAT device.

<https://docs.aws.amazon.com/vpn/latest/s2svpn/cgw-options.html>

upvoted 11 times

🗳️ 👤 **moonwalkeryj** Highly Voted 2 years, 4 months ago

If your customer gateway device is behind a network address translation (NAT) device, use the IP address of your NAT device.

<https://docs.aws.amazon.com/vpn/latest/s2svpn/cgw-options.html>

upvoted 6 times

🗳️ 👤 **seetpt** Most Recent 8 months, 3 weeks ago

Selected Answer: D

D is correct

upvoted 1 times

🗳️ 👤 **Mangesh_XI_mumbai** 1 year, 1 month ago

Selected Answer: D

D is correct.

upvoted 2 times

🗳️ 👤 **satamex** 1 year, 4 months ago

D absolutely

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: D

dddddddd

upvoted 2 times

🗳️ 👤 **Liongeek** 2 years, 1 month ago

Ans: D

upvoted 1 times

A company has a web application that is experiencing performance problems many times each night. A root cause analysis reveals sudden increases in CPU utilization that last 5 minutes on an Amazon EC2 Linux instance. A SysOps administrator must find the process ID (PID) of the service or process that is consuming more CPU.

What should the SysOps administrator do to collect the process utilization information with the LEAST amount of effort?

- A. Configure the Amazon CloudWatch agent procstat plugin to capture CPU process metrics.
- B. Configure an AWS Lambda function to run every minute to capture the PID and send a notification.
- C. Log in to the EC2 instance by using a .pem key each night. Then run the top command.
- D. Use the default Amazon CloudWatch CPU utilization metric to capture the PID in CloudWatch.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **moonwalkeryj** Highly Voted 2 years, 10 months ago

A

The procstat plugin enables you to collect metrics from individual processes. It is supported on Linux servers and on servers running Windows Server 2012 or later.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Agent-procstat-process-metrics.html>

upvoted 11 times

🗳️ **stallionaws** Most Recent 10 months ago

Selected Answer: A

The procstat plugin allows monitoring of specific processes based on criteria such as process name, command line, or user. This means you can capture CPU utilization data and PIDs for the processes causing high CPU usage.

upvoted 1 times

🗳️ **be9z** 1 year ago

Option D:

The default Amazon CloudWatch CPU utilization metric already provides aggregated data on CPU usage for the EC2 instance.

It requires no additional configuration or setup.

While it doesn't directly provide the specific process ID (PID), it's the simplest and least effort-intensive approach.

upvoted 1 times

🗳️ **flaacko** 9 months, 4 weeks ago

This is not true as the easiest way will be to use the procstat plugin by configuring on the CloudWatch agent. It will collect data and metrics for each individual process on your EC2 instance. So the correct option is A.

upvoted 1 times

🗳️ **Mangesh_XI_mumbai** 1 year, 7 months ago

Selected Answer: A

A is correct, procstat to be enabled.

upvoted 1 times

🗳️ **DZRomero** 1 year, 11 months ago

Opcion A

upvoted 1 times

🗳️ **anantk007** 1 year, 12 months ago

Selected Answer: A

Procstat plugin is the answer

upvoted 1 times

🗳️ **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaa

upvoted 1 times

🔍 👤 **Liongeek** 2 years, 7 months ago

ANS: A

upvoted 1 times

🔍 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: A

A. CloudWatch by default collects only aggregated CPU utilization, but can be configured with the Procstat plugin to collect more data.

upvoted 4 times

🔍 👤 **get_certified** 2 years, 9 months ago

Selected Answer: A

Procstat plugin allows to collect more data

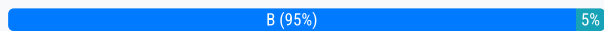
upvoted 2 times

A SysOps administrator configured AWS Backup to capture snapshots from a single Amazon EC2 instance that has one Amazon Elastic Block Store (Amazon EBS) volume attached. On the first snapshot, the EBS volume has 10 GiB of data. On the second snapshot, the EBS volume still contains 10 GiB of data, but 4 GiB have changed. On the third snapshot, 2 GiB of data have been added to the volume, for a total of 12 GiB. How much total storage is required to store these snapshots?

- A. 12 GiB
- B. 16 GiB
- C. 26 GiB
- D. 32 GiB

Suggested Answer: B

Community vote distribution



haxaffee Highly Voted 2 years, 4 months ago

Selected Answer: B

100% B. This scenario is explained at https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html#how_snapshots_work
upvoted 12 times

axdevops Highly Voted 1 year ago

1. First Snapshot: The entire volume is captured in the first snapshot, so it requires storage equivalent to the size of the volume, which is 10 GiB.
2. Second Snapshot: Only the changed blocks since the first snapshot are stored. In this case, 4 GiB have changed, so the storage required for the second snapshot is 4 GiB.
3. Third Snapshot: Only the changed blocks since the second snapshot are stored. In this case, 2 GiB have been added, so the storage required for the third snapshot is 2 GiB.

The total storage required is the sum of the storage for each snapshot:

$$10 \text{ GiB} + 4 \text{ GiB} + 2 \text{ GiB} = 16 \text{ GiB}$$

So, the total storage required to store these three snapshots is 16 GiB.

upvoted 11 times

peacefull Most Recent 1 month, 1 week ago

Selected Answer: B

AWS EBS snapshots are incremental: after the first full snapshot, each new snapshot stores only the data blocks that changed since the previous one, while unchanged blocks are merely referenced.

upvoted 1 times

baraadd 10 months, 1 week ago

Amazon EBS snapshots are incremental, meaning only the data changed after the first snapshot is saved. The first snapshot captures 10 GiB, the second snapshot only saves the 4 GiB that changed, and the third snapshot only saves the newly added 2 GiB. Therefore, the total storage required is the initial 10 GiB plus the total of 6 GiB for added/changed data (4 GiB + 2 GiB), making it 12 GiB in total, not 16 GiB.

upvoted 2 times

Mangesh_XI_mumbai 1 year, 1 month ago

Selected Answer: B

B is correct, $10 + 4 + 2 = 16 \text{ GiB}$

upvoted 3 times

fig 1 year, 1 month ago

Selected Answer: B

Answer is B: Original 10 + 4 (changed data) + 2 (new data) = 16

upvoted 3 times

🗨️ 👤 **csw23** 1 year, 3 months ago

this is a math question, not aws
upvoted 2 times

🗨️ 👤 **brtest** 1 year, 8 months ago

The first snapshot that you create from a volume is always a full snapshot. It includes all of the data blocks written to the volume at the time of creating the snapshot. Subsequent snapshots of the same volume are incremental snapshots. They include only changed and new data blocks written to the volume since the last snapshot was created.

upvoted 4 times

🗨️ 👤 **michaldavid** 2 years ago

Selected Answer: B

bbbbbbb

upvoted 2 times

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

ANS: B

upvoted 1 times

🗨️ 👤 **Capy** 2 years, 2 months ago

Selected Answer: B

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSSnapshots.html#how_snapshots_work

upvoted 2 times

🗨️ 👤 **Surferbolt** 2 years, 2 months ago

Selected Answer: D

D. Incremental snapshots.

upvoted 1 times

🗨️ 👤 **get_certified** 2 years, 3 months ago

Selected Answer: B

Due to incremental snapshots?

upvoted 2 times

A team is managing an AWS account that is a member of an organization in AWS Organizations. The organization has consolidated billing features enabled. The account hosts several applications.

A SysOps administrator has applied tags to resources within the account to reflect the environment. The team needs a report of the breakdown of charges by environment.

What should the SysOps administrator do to meet this requirement?

- A. Filter, map, and categorize resource groups in Tag Editor.
- B. Ensure that the organization's service control policies (SCPs) allow access to cost allocation tags.
- C. Ensure that the IAM credentials that are used to access Cost Explorer have permissions to group cost by tags.
- D. Activate the tag keys for cost allocation on the organization's management account.

Suggested Answer: A

Community vote distribution

D (100%)

🗳️ 👤 **moonwalkeryj** Highly Voted 2 years, 10 months ago

D

You must activate both types of tags separately before they can appear in Cost Explorer or on a cost allocation report.

<https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html>

upvoted 15 times

🗳️ 👤 **Surferbolt** Highly Voted 2 years, 8 months ago

Selected Answer: D

D. You need to activate the tags. It's not A because the question already said that resources has been tagged.

upvoted 10 times

🗳️ 👤 **jipark** 1 year, 10 months ago

activate keys !!

upvoted 1 times

🗳️ 👤 **flaacko** Most Recent 9 months, 4 weeks ago

Selected Answer: D

You can use cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the cost allocation tags to organize your resource costs on your cost allocation report, to make it easier for you to categorize and track your AWS costs.

upvoted 1 times

🗳️ 👤 **james2033** 1 year, 4 months ago

Selected Answer: D

To activate your tag keys

1) Sign in to the AWS Management Console and open the AWS Billing and Cost Management console at <https://console.aws.amazon.com/billing/>.

2) In the navigation pane, choose Cost allocation tags.

3) Select the tag keys that you want to activate.

4) Choose Activate.

See: <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/activating-tags.html>

upvoted 5 times

🗳️ 👤 **Mangesh_XI_mumbai** 1 year, 7 months ago

D is correct, activate tags. as per below link

upvoted 2 times

🗳️ 👤 **SebastianPlaiasu** 1 year, 12 months ago

A. Filter, map, and categorize resource groups in Tag Editor.

By using the Tag Editor, the administrator can filter resources based on specific tags, map those resources to resource groups, and then categorize them according to the desired environment. This will allow them to visualize and analyze the cost allocation based on the tagged environments.

Option D is incorrect because activating tag keys for cost allocation on the organization's management account is not a valid option. Tag activation is done on individual resources to enable cost allocation tagging, but it does not address the requirement of generating a breakdown of charges by environment.

upvoted 2 times

🗨️ 👤 **Niroljin** 2 years, 2 months ago

Selected Answer: D

same as the others. Tags were already used

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddd

upvoted 1 times

🗨️ 👤 **Daniel_Y** 2 years, 6 months ago

D is the answer

upvoted 1 times

🗨️ 👤 **pravinb** 2 years, 7 months ago

only D.

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Selected Answer: D

Ans: D

upvoted 1 times

🗨️ 👤 **hippius** 2 years, 8 months ago

Selected Answer: D

D is right answer

upvoted 1 times

🗨️ 👤 **get_certified** 2 years, 9 months ago

Selected Answer: D

Using cost allocation tags

upvoted 2 times

🗨️ 👤 **haxaffee** 2 years, 10 months ago

Selected Answer: D

Agree with moonwalkeryj on D

upvoted 1 times

A company uses an AWS CloudFormation template to provision an Amazon EC2 instance and an Amazon RDS DB instance. A SysOps administrator must update the template to ensure that the DB instance is created before the EC2 instance is launched. What should the SysOps administrator do to meet this requirement?

- A. Add a wait condition to the template. Update the EC2 instance user data script to send a signal after the EC2 instance is started.
- B. Add the DependsOn attribute to the EC2 instance resource, and provide the logical name of the RDS resource.
- C. Change the order of the resources in the template so that the RDS resource is listed before the EC2 instance resource.
- D. Create multiple templates. Use AWS CloudFormation StackSets to wait for one stack to complete before the second stack is created.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **RicardoD** Highly Voted 👍 3 years, 2 months ago
B is the answer

With the DependsOn attribute you can specify that the creation of a specific resource follows another. When you add a DependsOn attribute to a resource, that resource is created only after the creation of the resource specified in the DependsOn attribute.
upvoted 18 times

🗳️ 👤 **jkwek** Highly Voted 👍 3 years, 1 month ago
Answer is B.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-dependson.html>

Syntax
The DependsOn attribute can take a single string or list of strings.

"DependsOn" : [String, ...]

Example

The following template contains an AWS::EC2::Instance resource with a DependsOn attribute that specifies myDB, an AWS::RDS::DBInstance. When CloudFormation creates this stack, it first creates myDB, then creates Ec2Instance.
upvoted 8 times

🗳️ 👤 **jepwi** Most Recent 🕒 2 months, 2 weeks ago
Selected Answer: B

Depends on basically tells you that it will provision this resource right after the one that it depends on
upvoted 1 times

🗳️ 👤 **james2033** 9 months, 4 weeks ago
Selected Answer: B

Question's keyword: 'DB instances is created before the EC2 instance'.

Answer's keyword: DependsOn
upvoted 2 times

🗳️ 👤 **Mangesh_XI_mumbai** 1 year, 1 month ago
Selected Answer: B

dependsOn correct answer.
upvoted 1 times

🗳️ 👤 **michaldavid** 2 years ago
Selected Answer: B

bbbbbb
upvoted 1 times

🗳️ 👤 **Liongeek** 2 years, 1 month ago

Ans: B

upvoted 1 times

🗲️ 👤 **Surferbolt** 2 years, 2 months ago

Selected Answer: B

B. DependsOn attribute.

upvoted 1 times

🗲️ 👤 **Finger41** 2 years, 7 months ago

Selected Answer: B

B - <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-dependson.html>

upvoted 2 times

🗲️ 👤 **roka_ua** 2 years, 9 months ago

Selected Answer: B

Answer is B

upvoted 1 times

🗲️ 👤 **MrkJobs** 2 years, 12 months ago

Selected Answer: B

B is the answer:

"With the DependsOn attribute you can specify that the creation of a specific resource follows another. When you add a DependsOn attribute to a resource, that resource is created only after the creation of the resource specified in the DependsOn attribute."

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-dependson.html>

upvoted 2 times

🗲️ 👤 **szl0144** 2 years, 12 months ago

vote B

upvoted 1 times

🗲️ 👤 **ngthien041292** 3 years, 1 month ago

Selected Answer: B

Vote B

upvoted 2 times

A company hosts a static website on Amazon S3. The website is served by an Amazon CloudFront distribution with a default TTL of 86,400 seconds.

The company recently uploaded an updated version of the website to Amazon S3. However, users still see the old content when they refresh the site. A SysOps administrator must make the new version of the website visible to users as soon as possible.

Which solution meets these requirements?

- A. Adjust the TTL value for the DNS CNAME record that is pointing to the CloudFront distribution.
- B. Create an invalidation on the CloudFront distribution for the old S3 objects.
- C. Create a new CloudFront distribution. Update the DNS records to point to the new CloudFront distribution.
- D. Update the DNS record for the website to point to the S3 bucket.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **[Removed]** Highly Voted 1 year, 11 months ago

The correct answer is B. You can create an invalidation on the CloudFront distribution for the old S3 objects. This will remove the specified objects from CloudFront edge caches, forcing CloudFront to retrieve the updated content from the origin (in this case, the S3 bucket). This will make the new version of the website visible to users as soon as possible.

upvoted 8 times

🗳️ 👤 **jipark** 1 year, 10 months ago

thanks so much. "Invalidation is a to replicate the changes"

upvoted 2 times

🗳️ 👤 **OlehKom** Most Recent 7 months, 1 week ago

Selected Answer: B

Invalidation is a CloudFront feature that tells it to remove specific cached content immediately.

Once you create an invalidation for the old objects (like index.html), CloudFront will fetch the updated version from S3 and serve it to users.

upvoted 1 times

🗳️ 👤 **james2033** 1 year, 3 months ago

Selected Answer: B

86400 seconds = 1440 minutes = 24 hours = 1 day.

'Invalidating files using the console'

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Invalidation.html#Invalidation_Requests

upvoted 1 times

🗳️ 👤 **rajesh1232** 2 years, 1 month ago

can any one give exam topics sys ops all questions if any one subscribe recently,
please whats up me 9284423616 i will share some amount

upvoted 2 times

🗳️ 👤 **ambergmeh** 2 years, 1 month ago

TG ME @nomad_am

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbbbbb

upvoted 1 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

upvoted 1 times

🗳️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: B

B. Invalidate old versions so newer visits will cache the latest version.



upvoted 2 times

  **get_certified** 2 years, 9 months ago

Selected Answer: B

Invalidation is a to replicate the changes

upvoted 1 times

  **haxaffee** 2 years, 10 months ago

Selected Answer: B

B is correct. Invalidations can be created for the entire bucket content or specific paths.

upvoted 1 times

A SysOps administrator is responsible for managing a company's cloud infrastructure with AWS CloudFormation. The SysOps administrator needs to create a single resource that consists of multiple AWS services. The resource must support creation and deletion through the CloudFormation console.

Which CloudFormation resource type should the SysOps administrator create to meet these requirements?

- A. AWS::EC2::Instance with a cfn-init helper script
- B. AWS::OpsWorks::Instance
- C. AWS::SSM::Document
- D. Custom::MyCustomType

Suggested Answer: D

Community vote distribution

D (100%)

  **princajen**  1 year, 10 months ago

Selected Answer: D

Custom resources enable you to write custom provisioning logic in templates that AWS CloudFormation runs anytime you create, update (if you changed the custom resource), or delete stacks. For example, you might want to include resources that aren't available as AWS CloudFormation resource types. You can include those resources by using custom resources. That way you can still manage all your related resources in a single stack.




<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html>

upvoted 12 times

  **jipark** 10 months, 3 weeks ago

"For example, you might want to include resources that aren't available as AWS CloudFormation resource types."

upvoted 1 times

  **KikiNoviandi**  9 months, 2 weeks ago

Selected Answer: D

d correct

upvoted 1 times

  **Christina666** 11 months ago

Selected Answer: D

The Custom::MyCustomType is a CloudFormation custom resource type that allows you to define your own resource type using AWS Lambda-backed custom logic. This means you can use AWS Lambda functions to create and delete the custom resource.

By creating a custom resource backed by an AWS Lambda function, you have the flexibility to define and orchestrate multiple AWS services as a single resource within your CloudFormation stack. The custom Lambda function can handle the logic for creating and deleting the resources across multiple services in the desired order and configuration.

upvoted 2 times

  **Christina666** 11 months ago

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources-lambda.html>

upvoted 1 times

  **Christina666** 11 months ago

Option A (AWS::EC2::Instance with a cfn-init helper script) is used to create an EC2 instance, but it does not provide the ability to create multiple AWS services as a single resource. cfn-init is useful for bootstrapping EC2 instances with additional configuration and packages, but it's not designed to create other resources.

Option B (AWS::OpsWorks::Instance) is used to create an AWS OpsWorks instance, which is not applicable for the requirement of creating multiple AWS services as a single resource.

Option C (AWS::SSM::Document) is used to create an SSM (Systems Manager) document, which is used for automating tasks on instances using SSM, but it does not provide the functionality of creating multiple AWS services as a single resource.

upvoted 3 times

  **[Removed]** 11 months, 3 weeks ago

The correct answer is D. You can create a custom resource type, such as Custom::MyCustomType, to meet these requirements. A custom resource allows you to write custom provisioning logic in templates that AWS CloudFormation runs anytime you create, update, or delete stacks. You can use the custom resource to create a single resource that consists of multiple AWS services and supports creation and deletion through the CloudFormation console.

Received message. The correct answer is ****D****. You can create a custom resource type, such as `Custom::MyCustomType`, to meet these requirements. A custom resource allows you to write custom provisioning logic in templates that AWS CloudFormation runs anytime you create, update, or delete stacks. You can use the custom resource to create a single resource that consists of multiple AWS services and supports creation and deletion through the CloudFormation console.

upvoted 1 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddd

upvoted 2 times

A new website will run on Amazon EC2 instances behind an Application Load Balancer. Amazon Route 53 will be used to manage DNS records. What type of record should be set in Route 53 to point the website's apex domain name (for example, `company.com`) to the Application Load Balancer?

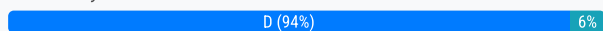
- A. CNAME
- B. SOA
- C. TXT
- D. ALIAS

Suggested Answer: D

Reference:

<https://docs.aws.amazon.com/govcloud-us/latest/UserGuide/setting-up-route53-zoneapex-elb.html>

Community vote distribution



[Removed] **Highly Voted** 1 year, 5 months ago

The correct answer is D. You should create an ALIAS record in Route 53 to point the website's apex domain name (e.g., company.com) to the Application Load Balancer. An ALIAS record is a Route 53-specific extension to DNS that allows you to map the apex domain name to another AWS resource, such as an Application Load Balancer.

upvoted 5 times

jipark 1 year, 4 months ago

thanks for explanation.

"Alias : apex domain name to another AWS resource, such as an Application Load Balancer."

upvoted 1 times

Will28Zhang28 **Most Recent** 11 months ago

Selected Answer: D

Not CNAME, as CNAME can not be applied to root domain level.

upvoted 3 times

KikiNoviandi 1 year, 3 months ago

Selected Answer: D

ALIA -> APEX

<https://docs.aws.amazon.com/govcloud-us/latest/UserGuide/setting-up-route53-zoneapex-elb.html>

upvoted 4 times

Gil80 1 year, 10 months ago

Selected Answer: D

Option D is the correct answer.

An ALIAS record is a Route 53 extension to DNS that provides a Route 53-specific alias for a resource record set. When Route 53 receives a DNS query for a domain name that is associated with an ALIAS record, Route 53 responds with the IP address of the resource that is specified in the record. An ALIAS record can be used to map an apex domain name to an Application Load Balancer. In this case, the ALIAS record points to the DNS name of the Application Load Balancer.

A CNAME record is a type of DNS record that points one domain name to another. However, CNAME records cannot be used to map an apex domain name to an Application Load Balancer.

An SOA (Start of Authority) record is a DNS record that specifies authoritative information about a DNS zone, such as the email address of the person responsible for managing the zone.

A TXT record is a type of DNS record that is used to associate some arbitrary text with a DNS record. It is not used to map domain names to IP addresses or to other domain names.

upvoted 4 times

🗨️ 👤 **michaldavid** 2 years ago

Selected Answer: D

dddddd

upvoted 1 times

🗨️ 👤 **Daniel_Y** 2 years ago

Only D, when pointing to resource names the record is Alias

upvoted 1 times

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

Ans: D

upvoted 2 times

🗨️ 👤 **Surferbolt** 2 years, 2 months ago

Selected Answer: D

D. SOA and TXT are definitely wrong. CNAME allows you to map names, but only for non zone apex addresses (ie m.example.com to mobile.example.com)

upvoted 3 times

🗨️ 👤 **neta1o** 2 years, 3 months ago

D, "Route 53 supports the alias resource record set, which lets you map your zone apex (e.g. example.com) DNS name to your load balancer DNS name."

upvoted 2 times

🗨️ 👤 **azure4life** 2 years, 3 months ago

Selected Answer: D

Definitely

upvoted 1 times

🗨️ 👤 **Chhotu_DBA** 2 years, 3 months ago

The answer is D

upvoted 2 times

🗨️ 👤 **ahmedtahalhosari** 2 years, 3 months ago

Selected Answer: A

We can't use ALIAS with root domain "apix"

upvoted 1 times

🗨️ 👤 **satamex** 1 year, 4 months ago

thats wrong.

upvoted 1 times

🗨️ 👤 **hiishii** 2 years, 4 months ago

I am going with D

upvoted 3 times

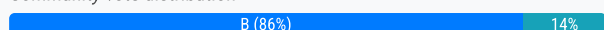
A company is implementing security and compliance by using AWS Trusted Advisor. The company's SysOps team is validating the list of Trusted Advisor checks that it can access.

Which factor will affect the quantity of available Trusted Advisor checks?

- A. Whether at least one Amazon EC2 instance is in the running state
- B. The AWS Support plan
- C. An AWS Organizations service control policy (SCP)
- D. Whether the AWS account root user has multi-factor authentication (MFA) enabled

Suggested Answer: B

Community vote distribution



🗳️ 👤 **Liongeek** Highly Voted 2 years, 7 months ago

Ans: B

Ref.: https://aws.amazon.com/premiumsupport/faqs/?nc1=h_ls

"Q: Which Trusted Advisor checks and features are available to all AWS customers?

AWS Basic Support and AWS Developer Support customers get access to 6 security checks (S3 Bucket Permissions, Security Groups - Specific Ports Unrestricted, IAM Use, MFA on Root Account, EBS Public Snapshots, RDS Public Snapshots) and 50 service limit checks. AWS Business Support, AWS Enterprise On-Ramp, and AWS Enterprise Support customers get access to all 115 Trusted Advisor checks (14 cost optimization, 17 security, 24 fault tolerance, 10 performance, and 50 service limits) and recommendations."

upvoted 15 times

🗳️ 👤 **OlehKom** Most Recent 7 months, 1 week ago

If you see Trusted Advisor - look for support plan key word, there is nothing you can do without it

upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 11 months ago

The factor that will affect the quantity of available Trusted Advisor checks is B, the AWS Support plan. The number of Trusted Advisor checks available to an account depends on the AWS Support plan that is associated with the account. Basic Support provides access to 7 core Trusted Advisor checks, Developer Support provides access to 29 core Trusted Advisor checks, while Business and Enterprise Support provide access to all Trusted Advisor checks and guidance.

upvoted 4 times

🗳️ 👤 **braveheart22** 2 years, 4 months ago

The correct answer is BBBB for sure.

upvoted 2 times

🗳️ 👤 **Pontimau** 2 years, 7 months ago

Selected Answer: D

For me, the answer it's D:

Whether the AWS account root user has multi-factor authentication (MFA) enabled

upvoted 1 times

🗳️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: B

B is the answer.

upvoted 3 times

🗳️ 👤 **haxaffee** 2 years, 10 months ago

Selected Answer: B

<https://aws.amazon.com/premiumsupport/plans/>

upvoted 3 times

A SysOps administrator is investigating issues on an Amazon RDS for MariaDB DB instance. The SysOps administrator wants to display the database load categorized by detailed wait events.

How can the SysOps administrator accomplish this goal?

- A. Create an Amazon CloudWatch dashboard.
- B. Enable Amazon RDS Performance Insights.
- C. Enable and configure Enhanced Monitoring.
- D. Review the database logs in Amazon CloudWatch Logs.

Suggested Answer: B

Reference:

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

Overview of the Performance Schema

The Performance Schema monitors server events. In this context, an event is a server action that consumes time. Performance Schema events are distinct from binlog events and scheduler events.

The PERFORMANCE_SCHEMA storage engine collects event data using instrumentation in the database source code. The engine stores collected events in tables in the performance_schema database. You can query performance_schema just as you can query any other tables. For more information, see [MySQL Performance Schema](#) in *MySQL Reference Manual*.

When the Performance Schema is enabled for Amazon RDS for MariaDB or MySQL, Performance Insights uses it to provide more detailed information. For example, Performance Insights displays DB load categorized by detailed wait events. You can use wait events to identify bottlenecks. Without the Performance Schema, Performance Insights reports user states such as inserting and sending, which don't help you identify bottlenecks.

Community vote distribution

B (100%)

CodePoet Highly Voted 3 years, 1 month ago

Selected Answer: B

Performance Insight is the right answer <https://aws.amazon.com/rds/performance-insights/>
upvoted 9 times

flaacko Most Recent 9 months, 4 weeks ago

Selected Answer: B

CPU utilization and database load can be related, but they are independent of each other. A database can be under high load, but have low CPU utilization. Likewise, a CPU could be experiencing high utilization levels, but be running only a few queries.

Performance Insights is a feature of Amazon RDS that provides visibility into the performance of SQL queries inside RDS databases.

upvoted 1 times

james2033 1 year, 3 months ago

Selected Answer: B

Amazon RDS Performance Insights
upvoted 1 times

TwinSpark 1 year, 8 months ago

Selected Answer: B

Correct answer B

"With the Performance Insights dashboard, you can visualize the database load on your Amazon RDS DB instance load and filter the load by waits, SQL statements, hosts, or users."

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

Instead Enhance monitoring Amazon RDS provides metrics in real time for the operating system (OS) that your DB instance runs on.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Monitoring.OS.overview.html

upvoted 4 times

🗲️ 👤 **Christina666** 1 year, 11 months ago

Selected Answer: B

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.UsingDashboard.AnalyzeDBLoad.html

DB load grouped by waits and top SQL queries is the default Performance Insights dashboard view. This combination typically provides the most insight into performance issues. DB load grouped by waits shows if there are any resource or concurrency bottlenecks in the database. In this case, the SQL tab of the top load items table shows which queries are driving that load.

upvoted 1 times

🗲️ 👤 **Christina666** 1 year, 11 months ago

Your typical workflow for diagnosing performance issues is as follows:

Review the Database load chart and see if there are any incidents of database load exceeding the Max CPU line.

If there is, look at the Database load chart and identify which wait state or states are primarily responsible.

Identify the digest queries causing the load by seeing which of the queries the SQL tab on the top load items table are contributing most to those wait states. You can identify these by the DB Load by Wait column.

Choose one of these digest queries in the SQL tab to expand it and see the child queries that it is composed of.

upvoted 1 times

🗲️ 👤 **[Removed]** 1 year, 11 months ago

To display the database load categorized by detailed wait events, the SysOps administrator should B, enable Amazon RDS Performance Insights. Performance Insights is an advanced performance monitoring feature that makes it easy to diagnose and solve performance problems on Amazon RDS. It displays the database load in an interactive graph, allowing you to analyze and troubleshoot the database workload. The load is categorized by SQL, waits, hosts, users, and other dimensions, providing detailed information about the sources of the load.

upvoted 1 times

🗲️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbbb

upvoted 1 times

🗲️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

upvoted 2 times

🗲️ 👤 **Atown** 2 years, 7 months ago

Selected Answer: B

B is the correct answers

upvoted 1 times

🗲️ 👤 **Finger41** 3 years, 1 month ago

Selected Answer: B

B - https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PerfInsights.Overview.ActiveSessions.html

upvoted 3 times

A company is planning to host an application on a set of Amazon EC2 instances that are distributed across multiple Availability Zones. The application must be able to scale to millions of requests each second.

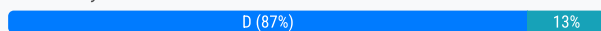
A SysOps administrator must design a solution to distribute the traffic to the EC2 instances. The solution must be optimized to handle sudden and volatile traffic patterns while using a single static IP address for each Availability Zone.

Which solution will meet these requirements?

- A. Amazon Simple Queue Service (Amazon SQS) queue
- B. Application Load Balancer
- C. AWS Global Accelerator
- D. Network Load Balancer

Suggested Answer: D

Community vote distribution



doggiecai Highly Voted 3 years, 2 months ago

Selected Answer: D

It's D.

"Network Load Balancer is optimized to handle sudden and volatile traffic patterns while using a single static IP address per Availability Zone."

<https://aws.amazon.com/elasticloadbalancing/network-load-balancer/>

upvoted 17 times

jipark 1 year, 10 months ago

static IP = NLB

upvoted 4 times

examaws Most Recent 6 months ago

Selected Answer: C

AWS Global Accelerator provides a static IP address that acts as a fixed entry point to your application, which can be beneficial for applications with sudden and volatile traffic patterns. It intelligently routes traffic to the optimal endpoint based on health, geography, and routing policies.

Application Load Balancer (B) is also a good option for distributing traffic but does not provide a single static IP address for each Availability Zone.

Network Load Balancer (D) can handle millions of requests per second and provides static IP addresses, but it is more suited for TCP/UDP traffic rather than HTTP/HTTPS.

Amazon SQS (A) is a message queuing service and does not directly distribute traffic to EC2 instances.

upvoted 1 times

OlehKom 7 months, 1 week ago

Selected Answer: D

Key words: High performance, static IP

upvoted 1 times

james2033 1 year, 3 months ago

Selected Answer: D

'Network Load Balancer' these words like introduction content at <https://aws.amazon.com/elasticloadbalancing/network-load-balancer/>

upvoted 1 times

Christina666 1 year, 11 months ago

Selected Answer: D

Option C (AWS Global Accelerator) can provide a single static anycast IP address for an application endpoint across multiple AWS regions. However, it does not provide a single static IP address per Availability Zone, which is specifically required in this scenario.

upvoted 2 times

[Removed] 1 year, 11 months ago

D. Network Load Balancer

upvoted 1 times

Gomer 2 years, 3 months ago

If the question was addressing EC2 instances across multiple regions, I think that Global Accelerator (in front of multiple network load balancers) would be the solution. However, it only specifies multiple AZ's, which IMHO implies the application is contained within a region. If it were global in scope, they'd need to specify that. I think the answer they want here is D because they are focusing primarily on performance requirement.

upvoted 2 times

🗳️ 👤 **Bhrino** 2 years, 3 months ago

Selected Answer: D

with this question you kinda just have to know what the services are used for. After reading this its between Application load balancer and Network load balancer. The reason its Network load balancer is because it says "Millions" and ALB cannot handle this much

upvoted 3 times

🗳️ 👤 **braveheart22** 2 years, 4 months ago

DDDD is the right way to go.

<https://aws.amazon.com/elasticloadbalancing/network-load-balancer/>

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

dddddddddd

upvoted 1 times

🗳️ 👤 **Daniel_Y** 2 years, 6 months ago

must be able to scale to millions of requests each second - the NLB can scale that much

upvoted 1 times

🗳️ 👤 **psdas** 2 years, 8 months ago

D, as it mentions static IP address.

upvoted 2 times

🗳️ 👤 **Surferbolt** 2 years, 8 months ago

D. Network Load balancer.

upvoted 1 times

🗳️ 👤 **Zulola** 2 years, 9 months ago

D..Network load Balancer

upvoted 1 times

🗳️ 👤 **221898** 3 years ago

Selected Answer: D

D. Network Load Balancer

upvoted 3 times

🗳️ 👤 **Seb** 3 years, 1 month ago

Selected Answer: C

Associate the static IP addresses provided by AWS Global Accelerator to regional AWS resources or endpoints, such as Network Load Balancers, Application Load Balancers, EC2 Instances, and Elastic IP addresses. The IP addresses are anycast from AWS edge locations so they provide onboarding to the AWS global network close to your users.

Easily move endpoints between Availability Zones or AWS Regions without needing to update your DNS configuration or change client-facing applications.

Dial traffic up or down for a specific AWS Region by configuring a traffic dial percentage for your endpoint groups. This is especially useful for testing performance and releasing updates.

Control the proportion of traffic directed to each endpoint within an endpoint group by assigning weights across the endpoints.

<https://aws.amazon.com/global-accelerator/faqs/>

upvoted 3 times

🗳️ 👤 **Rick365** 2 years, 9 months ago



it's D

upvoted 1 times

🗳️ 👤 **hexie** 2 years ago

Its D. Global Accelerator is not the recommended option for this scenario because it is designed for global application acceleration, focusing on directing traffic from users to the nearest AWS edge locations. It is not specifically designed for distributing traffic across EC2 instances within a single region or providing scalability and high availability within Availability Zones. In this case, Network Load Balancer (NLB) would be a more suitable choice :)

upvoted 1 times

  **altonh** 3 years, 1 month ago

Answer is D. It cannot be C because you only get 2 anycast IPs for the global accelerator.

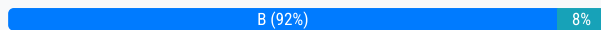
upvoted 4 times

A SysOps administrator is using AWS CloudFormation StackSets to create AWS resources in two AWS Regions in the same AWS account. A stack operation fails in one Region and returns the stack instance status of OUTDATED. What is the cause of this failure?

- A. The CloudFormation template changed on the local disk and has not been submitted to CloudFormation.
- B. The CloudFormation template is trying to create a global resource that is not unique.
- C. The stack has not yet been deployed to the Region.
- D. The SysOps administrator is using an old version of the CloudFormation API.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **haxaffee** Highly Voted 2 years, 10 months ago

Selected Answer: B

B as this is listed in <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-troubleshooting.html>
upvoted 12 times

🗳️ 👤 **Rick365** Highly Voted 2 years, 9 months ago

agree with haxaffee,

Common reasons for stack operation failure

Problem: A stack operation failed, and the stack instance status is OUTDATED.

Cause: There can be several common causes for stack operation failure.

The template could be trying to create global resources that must be unique but aren't, such as S3 buckets.

upvoted 5 times

🗳️ 👤 **jipark** 1 year, 10 months ago

not sure what this means : "reate global resources that must be unique but aren't"

upvoted 1 times

🗳️ 👤 **ahrentom** 1 year, 9 months ago

if Cloudformation want's to create a new S3 Bucket with a name but this name is always in use by another S3 Bucket. Bucket Names must be unique globally.

upvoted 2 times

🗳️ 👤 **examaws** Most Recent 6 months ago

Selected Answer: B

why not C. The stack has not yet been deployed to the Region.

Because : If the stack had not been deployed, the status would not be OUTDATED. Instead, it would indicate that the stack instance is in a state like NOT_CREATED or ROLLBACK_COMPLETE.

upvoted 1 times

🗳️ 👤 **silly_banana** 7 months, 2 weeks ago

Selected Answer: C

Ans is c

upvoted 2 times

🗳️ 👤 **ivaldesgo** 9 months, 4 weeks ago

Selected Answer: C

C. The stack has not yet been deployed to the Region.

upvoted 2 times

🗳️ 👤 **Aamee** 8 months ago

Read the question again... B is what the solution is for the issue discussed here.



upvoted 1 times

🗳️ 👤 **tgiv** 1 year, 2 months ago

Selected Answer: B

This can happen if CloudFormation tries to create a resource that already exists.

upvoted 1 times

  **tamng** 1 year, 6 months ago

B. The CloudFormation template is trying to create a global resource that is not unique.

upvoted 1 times

A SysOps administrator must configure Amazon S3 to host a simple nonproduction webpage. The SysOps administrator has created an empty S3 bucket from the

AWS Management Console. The S3 bucket has the default configuration in place.

Which combination of actions should the SysOps administrator take to complete this process? (Choose two.)

- A. Configure the S3 bucket by using the "Redirect requests for an object" functionality to point to the bucket root URL.
- B. Turn off the "Block all public access" setting. Allow public access by using a bucket ACL that contains <Permission>WEBSITE</Permission>.
- C. Turn off the "Block all public access" setting. Allow public access by using a bucket ACL that allows access to the AuthenticatedUsers grantee.
- D. Turn off the "Block all public access" setting. Set a bucket policy that allows "Principal": the s3:GetObject action.
- E. Create an index.html document. Configure static website hosting, and upload the index document to the S3 bucket.

Suggested Answer: DE

Community vote distribution

DE (100%)

  **princajen**  1 year, 10 months ago

Selected Answer: DE

Step 1: Create a bucket

Step 2: Enable static website hosting

Step 3: Edit Block Public Access settings

Step 4: Add a bucket policy that makes your bucket content publicly available

Step 5: Configure an index document

Step 6: Configure an error document

Step 7: Test your website endpoint

Step 8: Clean up

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/HostingWebsiteOnS3Setup.html>

upvoted 16 times

  **[Removed]**  11 months, 3 weeks ago

Correct Answer: DE



upvoted 1 times

  **s50600822** 1 year ago

Is there something missing: "Set a bucket policy that allows "Principal": the s3:GetObject action."

Like Principal wat ?

upvoted 4 times

  **eboehm** 11 months, 2 weeks ago

pretty sure its a typo and "the" should be *

upvoted 3 times



  **jipark** 10 months, 3 weeks ago

I also wonder it should be :

"Principal": "arn:aws:iam::111122223333:root",

"Action": "s3:PutObject"

upvoted 3 times

  **MrMLB** 1 year, 6 months ago

Selected Answer: DE

D & E.

upvoted 2 times

  **Liongeek** 1 year, 7 months ago

Selected Answer: DE

Ans: D&E

upvoted 2 times

A company is using an Amazon Aurora MySQL DB cluster that has point-in-time recovery, backtracking, and automatic backup enabled. A SysOps administrator needs to be able to roll back the DB cluster to a specific recovery point within the previous 72 hours. Restores must be completed in the same production DB cluster.

Which solution will meet these requirements?

- A. Create an Aurora Replica. Promote the replica to replace the primary DB instance.
- B. Create an AWS Lambda function to restore an automatic backup to the existing DB cluster.
- C. Use backtracking to rewind the existing DB cluster to the desired recovery point.
- D. Use point-in-time recovery to restore the existing DB cluster to the desired recovery point.

Suggested Answer: C

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/aurora-mysql-slow-snapshot-restore/>

Resolution

Note: If you receive errors when running AWS Command Line Interface (AWS CLI) commands, **make sure that you're using the most recent version of the AWS CLI.**

Amazon Aurora backs-up your cluster volume's changes automatically and continuously. The back-ups are retained for the length of your **backup retention period**. This continuous backup also means that you are able to restore your data to a new cluster, to any point in time within the retention period specified. This avoids the need for a lengthy binlog roll-forward process. Because you create a new cluster, there is no impact to performance or interruption to your original database.

When you initiate a clone, snapshot, or point in time restore, Amazon RDS calls the following APIs on your behalf:

- Either **RestoreDBClusterFromSnapshot** or **RestoreDBClusterToPointInTime**. This creates a new cluster and restores volume from Amazon Simple Storage Service (Amazon S3). This can take up to two hours to complete. This is because when you restore data to an Aurora cluster, all of the data must be brought in parallel from Amazon S3 to the six copies on your three AZs.
- **Cluster storage volume cloning** is a variation of **RestoreDBClusterToPointInTime**. It uses the copy-on-write protocol, and usually completes in a few minutes.

Community vote distribution

C (100%)

🗳️ 👤 **221898** **Highly Voted** 👍 1 year, 6 months ago

Selected Answer: C

C - <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>
upvoted 9 times

🗳️ 👤 **Finger41** **Highly Voted** 👍 1 year, 7 months ago

Selected Answer: C

C - <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>
upvoted 6 times

🗳️ 👤 **MrMLB** **Most Recent** 🕒 1 year ago

C for the win
upvoted 1 times

🗳️ 👤 **vjt** 1 year, 7 months ago

it is D. it is PITR and not backtracking. also Restored DB clusters are automatically associated with the default DB cluster and DB parameter groups.
Please clarify why C is suggested by people here?
upvoted 1 times

🗳️ 👤 **vjt** 1 year, 7 months ago

it is C. PITR will create a new DB cluster. however i am still confused with below aws info.

"You can restore a DB cluster to a specific point in time, creating a new DB cluster.

When you restore a DB cluster to a point in time, you can choose the default virtual private cloud (VPC) security group. Or you can apply a custom VPC security group to your DB cluster.

Restored DB clusters are automatically associated with the default DB cluster and DB parameter groups."

upvoted 4 times

  **Shriraj32** 1 year, 8 months ago

I think C is the answer.

upvoted 2 times

  **Mecdrox** 1 year, 8 months ago

Selected Answer: C

C. within 72 hours, backtracking is enabled. Must not restore to a new DB. C is the only solution that meets all requirements.

upvoted 5 times

  **dontcomplain** 1 year, 8 months ago

Please update answer should be C

upvoted 3 times

A user working in the Amazon EC2 console increased the size of an Amazon Elastic Block Store (Amazon EBS) volume attached to an Amazon EC2 Windows instance. The change is not reflected in the file system.
What should a SysOps administrator do to resolve this issue?

- A. Extend the file system with operating system-level tools to use the new storage capacity.
- B. Reattach the EBS volume to the EC2 instance.
- C. Reboot the EC2 instance that is attached to the EBS volume.
- D. Take a snapshot of the EBS volume. Replace the original volume with a volume that is created from the snapshot.

Suggested Answer: A

Community vote distribution

A (100%)

🗲️ 👤 **princajen** Highly Voted 👍 2 years, 4 months ago

Selected Answer: A

After you increase the size of an EBS volume, use the Windows Disk Management utility or PowerShell to extend the disk size to the new size of the volume. You can begin resizing the file system as soon as the volume enters the optimizing state.

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/recognize-expanded-volume-windows.html>

upvoted 15 times

🗲️ 👤 **tamng** Most Recent ⌚ 1 year ago

A. Extend the file system with operating system-level tools to use the new storage capacity.
upvoted 1 times

🗲️ 👤 **get_certified** 2 years, 3 months ago

Selected Answer: A

Extend the volume from Window Disk Manager tool
upvoted 2 times

A SysOps administrator is using Amazon EC2 instances to host an application. The SysOps administrator needs to grant permissions for the application to access an Amazon DynamoDB table. Which solution will meet this requirement?

- A. Create access keys to access the DynamoDB table. Assign the access keys to the EC2 instance profile.
- B. Create an EC2 key pair to access the DynamoDB table. Assign the key pair to the EC2 instance profile.
- C. Create an IAM user to access the DynamoDB table. Assign the IAM user to the EC2 instance profile.
- D. Create an IAM role to access the DynamoDB table. Assign the IAM role to the EC2 instance profile.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **princajen** Highly Voted 1 year, 10 months ago

Selected Answer: D

Access to Amazon DynamoDB requires credentials. Those credentials must have permissions to access AWS resources, such as an Amazon DynamoDB table or an Amazon Elastic Compute Cloud (Amazon EC2) instance. The following sections provide details on how you can use AWS Identity and Access Management (IAM) and DynamoDB to help secure access to your resources.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/authentication-and-access-control.html>

upvoted 8 times

🗳️ 👤 **Surferbolt** Highly Voted 1 year, 8 months ago

Selected Answer: D

D. From a security standpoint, always try to prioritize using roles over users.

upvoted 7 times

🗳️ 👤 **ad45** Most Recent 9 months ago

Selected Answer: D

Answer D : IAM role is a good practice, it could also be reused by another EC2 Instance, avoiding the multiplication of IAM Role

upvoted 2 times

🗳️ 👤 **airraid2010** 1 year, 1 month ago

Selected Answer: D

It is always safe to assign IAM role to a user.

upvoted 4 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddddd

upvoted 1 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

Ans: D

upvoted 1 times


A SysOps administrator wants to protect objects in an Amazon S3 bucket from accidental overwrite and deletion. Noncurrent objects must be kept for 90 days and then must be permanently deleted. Objects must reside within the same AWS Region as the original S3 bucket. Which solution meets these requirements?

- A. Create an Amazon Data Lifecycle Manager (Amazon DLM) lifecycle policy for the S3 bucket. Add a rule to the lifecycle policy to delete noncurrent objects after 90 days.
- B. Create an AWS Backup policy for the S3 bucket. Create a backup rule that includes a lifecycle to expire noncurrent objects after 90 days.
- C. Enable S3 Cross-Region Replication on the S3 bucket. Create an S3 Lifecycle policy for the bucket to expire noncurrent objects after 90 days.
- D. Enable S3 Versioning on the S3 bucket. Create an S3 Lifecycle policy for the bucket to expire noncurrent objects after 90 days.

Suggested Answer: D

Community vote distribution

D (100%)



  **kati2k22cz** Highly Voted 1 year, 10 months ago

Selected Answer: D

yes it's D.

<https://cloudacademy.com/blog/s3-lifecycle-policies-versioning-encryption-aws-security/>

upvoted 10 times

  **Anthony053** Highly Voted 1 year, 8 months ago

Enabling Versioning on the S3 bucket will protect accidental overwritten and deletion of data.

Versioning is same as 'Soft-delete' in an Azure!

upvoted 5 times

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

D is the right answer for keeping S3 objects updated

upvoted 2 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddddd

upvoted 2 times

A company has an application that customers use to search for records on a website. The application's data is stored in an Amazon Aurora DB cluster. The application's usage varies by season and by day of the week.

The website's popularity is increasing, and the website is experiencing slower performance because of increased load on the DB cluster during periods of peak activity. The application logs show that the performance issues occur when users are searching for information. The same search is rarely performed multiple times.

A SysOps administrator must improve the performance of the platform by using a solution that maximizes resource efficiency.

Which solution will meet these requirements?

- A. Deploy an Amazon ElastiCache for Redis cluster in front of the DB cluster. Modify the application to check the cache before the application issues new queries to the database. Add the results of any queries to the cache.
- B. Deploy an Aurora Replica for the DB cluster. Modify the application to use the reader endpoint for search operations. Use Aurora Auto Scaling to scale the number of replicas based on load.
- C. Use Provisioned IOPS on the storage volumes that support the DB cluster to improve performance sufficiently to support the peak load on the application.
- D. Increase the instance size in the DB cluster to a size that is sufficient to support the peak load on the application. Use Aurora Auto Scaling to scale the instance size based on load.

Suggested Answer: B

Community vote distribution

B (84%)

A (16%)

🗳️ 👤 **Surferbolt** Highly Voted 2 years, 8 months ago

Selected Answer: B

B. Caching will not solve the performance issue in this scenario, as the same search is rarely performed multiple times. Thus read replicas will be better.

upvoted 10 times

🗳️ 👤 **moonwalkeryj** Highly Voted 2 years, 10 months ago

B

It cannot be A, Because of this: The same search is rarely performed multiple times

upvoted 6 times

🗳️ 👤 **Pisces225** Most Recent 1 month, 2 weeks ago

Selected Answer: B

B - just adding to consensus.

upvoted 1 times

🗳️ 👤 **stallionaws** 10 months ago

Selected Answer: B

caching is not the option..rarely searched items..

upvoted 1 times

🗳️ 👤 **pekalyok** 1 year, 3 months ago

Selected Answer: A

implementing a caching layer with Amazon ElastiCache for Redis provides a scalable, cost-effective way to improve application performance by reducing database load, especially for read-heavy operations like searches. This approach addresses the core challenge of managing variable load efficiently.

upvoted 1 times

🗳️ 👤 **flaacko** 9 months, 4 weeks ago

A is wrong because there is no need for caching as the same search is rarely performed multiple times.

upvoted 2 times

🗳️ 👤 **Mangesh_XI_mumbai** 1 year, 7 months ago

Selected Answer: B

keyword : read replica

upvoted 2 times

🗨️ 👤 **pekalyok** 1 year, 3 months ago

wrong: here is why: "While adding Aurora Replicas and using the reader endpoint for search operations can distribute the read load and improve performance, this solution primarily increases database capacity rather than reducing unnecessary loads. It's more cost-intensive and may still face bottlenecks during unprecedented spikes in demand, as the underlying issue of inefficient query processing during peak times isn't directly addressed."

upvoted 1 times

🗨️ 👤 **callspace** 1 year, 9 months ago

Selected Answer: B

B It is. A is not possible because question clearly says (The same search is rarely performed multiple times.)

upvoted 2 times

🗨️ 👤 **braveheart22** 2 years, 4 months ago

BBBBB is the right answer in this case because caching with Elasticache for Redis will not help resolve the issue since the same search is rarely performed multiple times.

upvoted 5 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbbbbb

upvoted 3 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

Same search is rarely performed so Elasticache isn't needed

upvoted 4 times

🗨️ 👤 **get_certified** 2 years, 9 months ago

Selected Answer: B

Read replicas will solve the problem.

upvoted 3 times

🗨️ 👤 **princajen** 2 years, 9 months ago

Selected Answer: B

I'm voting B!

upvoted 2 times

🗨️ 👤 **Rick365** 2 years, 9 months ago

Selected Answer: B

The same search is rarely performed multiple times. Read Replicas will solve this issue

upvoted 1 times

🗨️ 👤 **CiCa560** 2 years, 10 months ago

Selected Answer: A

A - The question specifically talks about performance issues - ElasticCache will resolve the performance issue.

upvoted 3 times

🗨️ 👤 **wookchan** 1 year, 11 months ago

It's B, because of "The same search is rarely performed multiple times.

upvoted 2 times

🗨️ 👤 **pekalyok** 1 year, 3 months ago

its A cuz implementing a caching layer with Amazon ElastiCache for Redis provides a scalable, cost-effective way to improve application performance by reducing database load, especially for read-heavy operations like searches. This approach addresses the core challenge of managing variable load efficiently.

upvoted 1 times

🗨️ 👤 **kati2k22cz** 2 years, 10 months ago

Selected Answer: B

B. Here the AWS document

https://docs.amazonaws.cn/en_us/AmazonRDS/latest/AuroraUserGuide/aurora-replicas-adding.html

upvoted 1 times

A company uses AWS Organizations to manage multiple AWS accounts. Corporate policy mandates that only specific AWS Regions can be used to store and process customer data. A SysOps administrator must prevent the provisioning of Amazon EC2 instances in unauthorized Regions by anyone in the company.



What is the MOST operationally efficient solution that meets these requirements?

- A. Configure AWS CloudTrail in all Regions to record all API activity. Create an Amazon EventBridge (Amazon CloudWatch Events) rule in all unauthorized Regions for ec2:RunInstances events. Use AWS Lambda to terminate the launched EC2 instances.
- B. In each AWS account, create a managed IAM policy that uses a Region condition to deny the ec2:RunInstances action in all unauthorized Regions. Attach this policy to all IAM groups in each AWS account.
- C. In each AWS account, create an IAM permissions boundary policy that uses a Region condition to deny the ec2:RunInstances action in all unauthorized Regions. Attach the permissions boundary policy to all IAM users in each AWS account.
- D. Create a service control policy (SCP) in AWS Organizations to deny the ec2:RunInstances action in all unauthorized Regions. Attach this policy to the root level of the organization.

Suggested Answer: D

Community vote distribution

D (100%)

  **haxaffee** Highly Voted 2 years, 4 months ago

Selected Answer: D



Def D here. If this question has something like "Use Control Tower and setup the landing zone with provisioned regions" that would take be correct but there is no Control Tower option on this one.

upvoted 10 times

  **Mangesh_XI_mumbai** 1 year, 1 month ago

you are right if control tower is mentioned somewhere it would be easy to answer this, that is what they tricked in the question here, coz without control tower too, we can use SCPs using organization. otherwise most of them would go for A.



upvoted 1 times

  **jipark** 1 year, 4 months ago

chatGPT says :

"Option D provides the most efficient solution because it offers centralized control, granularity, and automatic enforcement across the organization's accounts, without the need for implementing and maintaining additional mechanisms like CloudTrail, CloudWatch Events, Lambda functions, IAM policies, or permissions boundaries."

upvoted 1 times

  **kolysh** 1 year, 5 months ago

What exact control tower you are talking about?

upvoted 1 times

  **Pisces225** Most Recent 1 month, 2 weeks ago

Selected Answer: D

D - just adding to consensus.

upvoted 1 times

  **tamng** 1 year ago

D. Create a service control policy (SCP) in AWS Organizations to deny the ec2:RunInstances action in all unauthorized Regions. Attach this policy to the root level of the organization.

upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: D

ddddddddd

upvoted 2 times

  **Liongeek** 2 years, 1 month ago

Ans: D

upvoted 1 times

A company's public website is hosted in an Amazon S3 bucket in the us-east-1 Region behind an Amazon CloudFront distribution. The company wants to ensure that the website is protected from DDoS attacks. A SysOps administrator needs to deploy a solution that gives the company the ability to maintain control over the rate limit at which DDoS protections are applied.

Which solution will meet these requirements?



- A. Deploy a global-scoped AWS WAF web ACL with an allow default action. Configure an AWS WAF rate-based rule to block matching traffic. Associate the web ACL with the CloudFront distribution.
- B. Deploy an AWS WAF web ACL with an allow default action in us-east-1. Configure an AWS WAF rate-based rule to block matching traffic. Associate the web ACL with the S3 bucket.
- C. Deploy a global-scoped AWS WAF web ACL with a block default action. Configure an AWS WAF rate-based rule to allow matching traffic. Associate the web ACL with the CloudFront distribution.
- D. Deploy an AWS WAF web ACL with a block default action in us-east-1. Configure an AWS WAF rate-based rule to allow matching traffic. Associate the web ACL with the S3 bucket.

Suggested Answer: A

Community vote distribution

A (80%)

C (20%)

  **kati2k22cz** Highly Voted 2 years, 10 months ago

Selected Answer: A

yes, A is the answer

<https://docs.aws.amazon.com/waf/latest/developerguide/classic-web-acl-default-action.html>

upvoted 10 times

  **Aamee** Most Recent 8 months ago

Selected Answer: A

A is the correct ans. here

upvoted 1 times

  **ignajtpolandstrong** 1 year, 4 months ago

Selected Answer: A

Your configuration of your own rules and rule groups depends in part on whether you want to allow or block most web requests. For example, if you want to allow most requests, you would set the web ACL default action to Allow, and then add rules that identify web requests that you want to block, such as the following:

Requests that originate from IP addresses that are making an unreasonable number of requests

Requests that originate from countries that either you don't do business in or are the frequent source of attacks

Requests that include fake values in the User-agent header

Requests that appear to include malicious SQL code

<https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-default-action.html>

upvoted 4 times

  **tamng** 1 year, 6 months ago

A is Correct


upvoted 1 times

  **axdevops** 1 year, 6 months ago

Answer: A

<https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-default-action.html>

upvoted 1 times

  **teo2157** 1 year, 7 months ago

I think the key point here is "control over the rate limit at which DDoS protections are applied.", if you want to maintain this control A is the answer, the other option C blocks by default everything.

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 1 month ago

Selected Answer: C

With rate-based rules, you only have Block, Count and Captcha. You don't have Allow. So A is incorrect. C is the good one.

upvoted 3 times

🗨️ 👤 **[Removed]** 2 years, 1 month ago

My mistake, answer A is correct, my mistake, please delete the messages.

upvoted 5 times

🗨️ 👤 **jipark** 1 year, 10 months ago

thanks for clarification !!

CloudFront is global -> need Global WAF

rate base rule can only block -> default allow

upvoted 3 times

🗨️ 👤 **[Removed]** 2 years, 1 month ago

With rate-based rules, you only have Block, Count and Captcha. You don't have Allow. So A is incorrect. C is the good one.

upvoted 1 times

🗨️ 👤 **brtest** 2 years, 2 months ago

I think is A because to deliver content to users with lower latency, Amazon CloudFront uses a global network of edge locations (edge locations and regional edge caches) around the world

upvoted 2 times

🗨️ 👤 **CVDON** 2 years, 4 months ago

Its global because you your are using cloudfront and it has to protect all the edge locations

upvoted 2 times

🗨️ 👤 **Spike2020** 2 years, 5 months ago

why not C? default action is blocked but matching traffic allowed. That means only us region is allowed with a rate limit.

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: A

I think A

upvoted 1 times

🗨️ 👤 **Vanfox** 2 years, 8 months ago

Why global scoped? Answer should be B.

upvoted 1 times

🗨️ 👤 **Vanfox** 2 years, 8 months ago

After reviewing I see it is A due to Cloudfront Distribution.

upvoted 7 times

A SysOps administrator developed a Python script that uses the AWS SDK to conduct several maintenance tasks. The script needs to run automatically every night.

What is the MOST operationally efficient solution that meets this requirement?

- A. Convert the Python script to an AWS Lambda function. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the function every night.
- B. Convert the Python script to an AWS Lambda function. Use AWS CloudTrail to invoke the function every night.
- C. Deploy the Python script to an Amazon EC2 instance. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule the instance to start and stop every night.
- D. Deploy the Python script to an Amazon EC2 instance. Use AWS Systems Manager to schedule the instance to start and stop every night.


Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/events/RunLambdaSchedule.html>

Community vote distribution

A (100%)

 **airraid2010** Highly Voted 2 years, 1 month ago

Selected Answer: A

Converting the Python script to an AWS Lambda function and using an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the function every night is a serverless and fully managed solution, which means that the administrator does not have to worry about managing the underlying infrastructure or scaling.

AWS Lambda allows for easy deployment, scaling, and management of code in response to events. EventBridge (CloudWatch Events) provides a reliable way to schedule events and to invoke AWS services on a schedule.

upvoted 7 times

 **Aamee** Most Recent 8 months ago

Selected Answer: A

Def. A as it's asks for 'Operationally efficient' method.

upvoted 1 times


 **jipark** 1 year, 10 months ago

Selected Answer: A

the reason for not C : "schedule the instance to start and stop every night."

-> restart instance every day sounds non-sense.

upvoted 1 times

 **Gomer** 2 years, 3 months ago

Everyone seems to choose A. without a second thought, and there is a default metric for Lambda invocation errors. However, this doesn't deal with the alarms that Lambda may generate while running. Dealing with this requires CloudWatch logs to use patterns such as ?ERROR to parse the log files and generate an alert. Since the question doesn't seem that exotic or specific, "A." is probably the right choice, but this is a tricky question. If you think it isn't, read: <https://aws.amazon.com/blogs/mt/get-notified-specific-lambda-function-error-patterns-using-cloudwatch/>

upvoted 1 times

 **Gomer** 2 years, 3 months ago

SORRY, DISREGARD MY COMMENT. I MEANT IT FOR QUESTION 37.

upvoted 2 times

 **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaaaa

upvoted 2 times

 **Liongeek** 2 years, 7 months ago

Selected Answer: A

Ans: A

upvoted 1 times

  **Surferbolt** 2 years, 8 months ago

Selected Answer: A

A. Lambda is good enough.



upvoted 1 times

  **get_certified** 2 years, 9 months ago

Selected Answer: A

Here a lambda function will be the most efficient solution.

upvoted 2 times

  **CiCa560** 2 years, 10 months ago

Selected Answer: A

Question asks for 'the MOST operationally efficient solution'

Deploying a Python script to an EC2 instance for execution would not be efficient

upvoted 1 times

A SysOps administrator must create a solution that immediately notifies software developers if an AWS Lambda function experiences an error. Which solution will meet this requirement?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic with an email subscription for each developer. Create an Amazon CloudWatch alarm by using the Errors metric and the Lambda function name as a dimension. Configure the alarm to send a notification to the SNS topic when the alarm state reaches ALARM.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic with a mobile subscription for each developer. Create an Amazon EventBridge (Amazon CloudWatch Events) alarm by using the LambdaError as the event pattern and the SNS topic name as a resource. Configure the alarm to send a notification to the SNS topic when the alarm state reaches ALARM.
- C. Verify each developer email address in Amazon Simple Email Service (Amazon SES). Create an Amazon CloudWatch rule by using the LambdaError metric and developer email addresses as dimensions. Configure the rule to send an email through Amazon SES when the rule state reaches ALARM.
- D. Verify each developer mobile phone in Amazon Simple Email Service (Amazon SES). Create an Amazon EventBridge (Amazon CloudWatch Events) rule by using Error as the event pattern and the Lambda function name as a resource. Configure the rule to send a push notification through Amazon SES when the rule state reaches ALARM.

Suggested Answer: A

Community vote distribution

A (100%)

ovilla **Highly Voted** 2 years, 8 months ago

It's A

upvoted 9 times

by116549 **Highly Voted** 2 years, 4 months ago

Not C or D as SES is used for marketing emails.

Leaves us with A or B.

With B AWS config, Cloudtrail or Cloudwatch can be the source no a SNS topic.

A seems the best reasonable solution

upvoted 8 times

willows **Most Recent** 8 months ago

A.

<https://aws.amazon.com/blogs/mt/get-notified-specific-lambda-function-error-patterns-using-cloudwatch/>

upvoted 1 times

tamng 1 year ago

A is Correct

upvoted 1 times

jipark 1 year, 4 months ago

Selected Answer: A

the reason for not B :

"SNS topic name as a resource" (X) vs "Lambda function name as a dimension" (O)

-> cloud watch set Lambda name as dimension.

upvoted 2 times

Gomer 1 year, 9 months ago

Everyone seems to choose A. without a second thought, and there is a default metric for Lambda invocation errors. However, this doesn't deal with the alarms that Lambda may generate while running. Dealing with this requires CloudWatch logs to use patterns such as ?ERROR to parse the log files and generate an alert. Since the question doesn't seem that exotic or specific, "A." is probably the right choice, but this is a tricky question. If you think it isn't, read: <https://aws.amazon.com/blogs/mt/get-notified-specific-lambda-function-error-patterns-using-cloudwatch/>

upvoted 2 times

  **eboehm** 1 year, 5 months ago

the question doesnt ask how to deal with the error/alarm. It just mentioned that someone needs notified when an error occurs. Once notified they can view the logs and resolve the error. Now ideally, yes, it should be a log subscription which would notify and send details about what caused the error, but the question did not ask for this

upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: A

aaaaaaa


upvoted 1 times

  **Liongeek** 2 years, 1 month ago

Selected Answer: A

Ans: A



upvoted 1 times

  **Finger41** 2 years, 7 months ago

Selected Answer: A

A - <https://aws.amazon.com/blogs/mt/get-notified-specific-lambda-function-error-patterns-using-cloudwatch/>

upvoted 5 times

  **psou7** 2 years, 8 months ago

Selected Answer: A

A is correct

upvoted 4 times

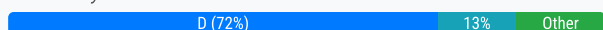
A company has a private Amazon S3 bucket that contains sensitive information. A SysOps administrator needs to keep logs of the IP addresses from authentication failures that result from attempts to access objects in the bucket. The logs must be stored so that they cannot be overwritten or deleted for 90 days.

Which solution will meet these requirements?

- A. Create an AWS CloudTrail trail. Configure the log files to be saved to Amazon CloudWatch Logs. Configure the log group with a retention period of 90 days.
- B. Create an AWS CloudTrail trail. Configure the log files to be saved to a different S3 bucket. Turn on CloudTrail log file integrity validation for 90 days.
- C. Turn on access logging for the S3 bucket. Configure the access logs to be saved to Amazon CloudWatch Logs. Configure the log group with a retention period of 90 days.
- D. Turn on access logging for the S3 bucket. Configure the access logs to be saved in a second S3 bucket. Turn on S3 Object Lock on the second S3 bucket, and configure a default retention period of 90 days.

Suggested Answer: D

Community vote distribution



kati2k22cz Highly Voted 2 years, 10 months ago

Selected Answer: D

D.

Learn more here: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html>

upvoted 11 times

princajen Highly Voted 2 years, 10 months ago

I would have thought is A, but after reading I found this: "CloudTrail does not deliver logs for requests that fail authentication (in which the provided credentials are not valid). However, it does include logs for requests in which authorization fails (AccessDenied) and requests that are made by anonymous users."

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html>

upvoted 10 times

cosmogen 2 years, 7 months ago

You are right, but read again the question: "keep logs of the IP addresses from authentication failures" that "result from attempts to access objects in the bucket". "That result from attempts to access objects in the bucket", for me it's mean authorization fails (AccessDenied). So, for me B, C, D technically are not possible. Vote for A

upvoted 3 times

Dinya_jui Most Recent 5 months, 4 weeks ago

Selected Answer: D

I would have thought is A, but after reading I found this: "CloudTrail does not deliver logs for requests that fail authentication (in which the provided credentials are not valid). However, it does include logs for requests in which authorization fails (AccessDenied) and requests that are made by anonymous users."

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html>

upvoted 1 times

numark 7 months, 1 week ago

D: S3 Access Logging allows you to capture details of requests made to your S3 bucket, including failed attempts. This is critical for tracking authentication failures. By configuring the access logs to be saved to a second S3 bucket, you can separate the logs from the sensitive data in the original bucket, adding an extra layer of security and compliance. S3 Object Lock can be used to prevent object versions from being deleted or overwritten for a specified retention period. By turning on S3 Object Lock and setting the retention period to 90 days, you ensure that the access logs are immutable for the required duration. S3 Object Lock enforces a Write Once, Read Many (WORM) model, which is ideal for compliance and security use cases.

upvoted 1 times

joshnort 1 year, 3 months ago

Selected Answer: D

[https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html#:~:text=CloudTrail%20does%20not%20deliver%20logs%20for%20requests%20that%20fail%20authentication%20\(in%20which%20the%20provided%20](https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html#:~:text=CloudTrail%20does%20not%20deliver%20logs%20for%20requests%20that%20fail%20authentication%20(in%20which%20the%20provided%20)
upvoted 4 times

Selected Answer: D

upvoted 3 times

Selected Answer: D

upvoted 5 times

Selected Answer: D

upvoted 2 times

upvoted 1 times

Selected Answer: D

upvoted 2 times

Selected Answer: B

upvoted 1 times

upvoted 1 times

Selected Answer: D

upvoted 8 times

Selected Answer: B

upvoted 1 times

🗨️ 👤 **Gomer** 2 years, 1 month ago

Integrity validation doesn't stop someone from changing/deleting logs, it just detects it. The real solution requires protection of the files through versioning or object lock. After wasting a day analyzing all the solutions, I found real problems with each one as I listed. That is why the voting is fairly balanced. Something is wrong with the responses given

upvoted 1 times

🗨️ 👤 **Gomer** 2 years, 1 month ago

NOT A: Can't configure CloudTrail to store logs in CloudWatch Logs. CloudTrail uses S3 bucket. CloudWatch Logs is not applicable.
<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/get-and-view-cloudtrail-log-files.html>

NOT B: "Integrity validation" is only designed to detect changes or deletions of CloudTrail logs. It depends on other security measures to block this.
<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

NOT C: Server access logging only delivers access logs for a source bucket to a target bucket. CloudWatch log group is not applicable.
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/enable-server-access-logging.htm>

NOT D: "S3 buckets with S3 Object Lock can't be used as destination buckets for server access logs."
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html>
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/enable-server-access-logging.html>

upvoted 4 times

🗨️ 👤 **Gomer** 2 years, 1 month ago

Both "A" and "C" are clearly not possible, and are excluded (IMHO). However, "B" and "D" also appear to be excluded for reasons cited. I do lean towards "D" because one clear requirement is to block log file deletion, not just detect it or automate it after 90 days.

IMHO the real solution (not listed) would create a CloudTrail trail that logs S3 Data Events in separate bucket, enable S3 Object Lock on that bucket with a retention period of 90 days, and enable Integrity Validation to detect any possible changes/deletions. Then I'd also figure out a lifecycle policy or some method to delete the logs sometimes after the 90 day requirement.

upvoted 1 times

🗨️ 👤 **vhernan** 2 years, 2 months ago

Selected Answer: B

By creating a CloudTrail trail, you can log all API calls made to the S3 bucket, including authentication failures. The logs can be saved to a separate S3 bucket to isolate them from the main bucket and provide an additional layer of security. Turning on CloudTrail log file integrity validation ensures that the logs cannot be modified or deleted without detection. The retention period for the logs can be set to 90 days to meet the requirements specified in the question.

upvoted 1 times

🗨️ 👤 **braveheart22** 2 years, 4 months ago

This question is really tricky, but after reading the question very carefully, I will definitely go with BBBBBB.

upvoted 4 times

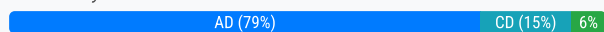
A SysOps administrator migrates NAT instances to NAT gateways. After the migration, an application that is hosted on Amazon EC2 instances in a private subnet cannot access the internet.

Which of the following are possible reasons for this problem? (Choose two.)

- A. The application is using a protocol that the NAT gateway does not support.
- B. The NAT gateway is not in a security group.
- C. The NAT gateway is in an unsupported Availability Zone.
- D. The NAT gateway is not in the Available state.
- E. The port forwarding settings do not allow access to internal services from the internet.

Suggested Answer: CD

Community vote distribution



piavik Highly Voted 2 years, 2 months ago

Selected Answer: AD

A - NAT gateway does not support IPv6
D - NAT gateway itself has issues or not yet ready
upvoted 7 times

griggrig Most Recent 11 months, 1 week ago

A and D

Causes

The cause of this problem might be one of the following:

The NAT gateway is not ready to serve traffic.

Your route tables are not configured correctly.

Your security groups or network ACLs are blocking inbound or outbound traffic.

You're using an unsupported protocol.

upvoted 1 times

willows 1 year, 1 month ago

A - TCP, UDP & ICMP only

D - Must be in available state

<https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-troubleshooting.html#nat-gateway-troubleshooting-no-internet-connection>

upvoted 2 times

ExamGuru727 1 year, 3 months ago

Selected Answer: CD

<https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-troubleshooting.html#nat-gateway-troubleshooting-unsupported-az>

From the doc above:

- AZ is unsupported - C

- Must be in Available state - D

upvoted 1 times

kret 1 year, 3 months ago

Your link points to article to troubleshoot NATGW creation errors. Question clearly says that NATGW was created, but is not working. So it can't be Unsupported AZ.

upvoted 1 times

🗄️ 👤 **Learning4life** 1 year, 5 months ago

Selected Answer: AD

<https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-troubleshooting.html#nat-gateway-troubleshooting-no-internet-connection>
upvoted 2 times

🗄️ 👤 **tamng** 1 year, 6 months ago

A D are correct answers
upvoted 1 times

🗄️ 👤 **tamng** 1 year, 6 months ago

Link: <https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-troubleshooting.html>
upvoted 1 times

🗄️ 👤 **callspace** 1 year, 9 months ago

Selected Answer: CD

Only CD Security group, port forwarding and protocol do not play a direct role in accessing the internet. As EC2 instances are in private subnet, 0.0.0.0/0 - nat-gateway-id entry is required to access internet and that option is not mentioned in the answers so CD are left as possible answers.
upvoted 2 times

🗄️ 👤 **Niroljin** 2 years, 2 months ago

Selected Answer: AD

A & D
Wrong protocol and nat not available
upvoted 2 times

🗄️ 👤 **caputmundi666** 2 years, 3 months ago

Selected Answer: AD

Please correct answer to AD
upvoted 2 times

🗄️ 👤 **caputmundi666** 2 years, 3 months ago

Please correct answer to AD
upvoted 1 times

🗄️ 👤 **Gomer** 2 years, 3 months ago

Selected Answer: AD

A & D for now. My research suggests that either A, C or D could be correct (see quotes and URL below). However, in analyzing the wording, I think "C" might be incorrect IF the NAT Gateway creation fails due to NotAvailableInZone error. If the NAT Gateway "object" still get's created, but is just not available due to the error, then C is back in play. Not sure if question is tricky or is just bad.
A: "Ensure that your connection is using a TCP, UDP, or ICMP protocol only."
C: If creation of NAT Gateway generates a "NotAvailableInZone" error, the "Availability Zone is unsupported" because it is constrained from being expanded
D: "Check that the NAT gateway is in the Available state." and "here may have been an error when it was created") (see: "Troubleshoot NAT gateways")
<https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-troubleshooting.html>
upvoted 4 times

🗄️ 👤 **jipark** 1 year, 10 months ago

I agree..
upvoted 1 times

🗄️ 👤 **FKZ** 2 years, 4 months ago

Selected Answer: AE

For sure.
upvoted 2 times

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

A and E

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html>
upvoted 1 times

🗄️ 👤 **sualej** 2 years, 5 months ago

A y D, as documentation said.
The cause of this problem might be one of the following:

The NAT gateway is not ready to serve traffic.
Your route tables are not configured correctly.
Your security groups or network ACLs are blocking inbound or outbound traffic.
You're using an unsupported protocol.

In case Availability Zone you were not able to create the NAT Gateway as you were got an error.

upvoted 1 times

🗨️ 👤 **BietTuot** 2 years, 6 months ago

Selected Answer: AD

Instances cannot access the internet

Problem

You created a public NAT gateway and followed the steps to test it, but the ping command fails, or your instances in the private subnet cannot access the internet.

Causes

The cause of this problem might be one of the following:

The NAT gateway is not ready to serve traffic.

Your route tables are not configured correctly.

Your security groups or network ACLs are blocking inbound or outbound traffic.

You're using an unsupported protocol.

Solution

Check the following information:

Check that the NAT gateway is in the Available state. In the Amazon VPC console, go to the NAT Gateways page and view the status information in the details pane. If the NAT gateway is in a failed state, there may have been an error when it was created. For more information, see NAT gateway creation fails.

Reference: <https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-troubleshooting.html>

upvoted 3 times

🗨️ 👤 **MrMLB** 2 years, 6 months ago

The possible reasons for this problem are options A and D: the application is using a protocol that the NAT gateway does not support, or the NAT gateway is not in the Available state. If the application is using a protocol that the NAT gateway does not support, it will not be able to access the internet through the NAT gateway. Similarly, if the NAT gateway is not in the Available state, it will not be able to provide internet access to the EC2 instances. The other options, B, C, and E, are not possible reasons for this problem.

upvoted 1 times

🗨️ 👤 **mautpo** 2 years, 7 months ago

Selected Answer: CD

Thinking about A, I believe its not a matter of protocol once application was behind NAT Instances and I was supposed to be working...Thats why I vote for C and D.

upvoted 2 times

🗨️ 👤 **piavik** 2 years, 2 months ago

NAT gateway does not support IPV6, so A is correct

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: D

Ref.<https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-troubleshooting.html#nat-gateway-troubleshooting-no-internet-connection>

upvoted 2 times

A company runs an application on an Amazon EC2 instance. A SysOps administrator creates an Auto Scaling group and an Application Load Balancer (ALB) to handle an increase in demand. However, the EC2 instances are failing the health check. What should the SysOps administrator do to troubleshoot this issue?




- A. Verify that the Auto Scaling group is configured to use all AWS Regions.
- B. Verify that the application is running on the protocol and the port that the listener is expecting.
- C. Verify the listener priority in the ALB. Change the priority if necessary.
- D. Verify the maximum number of instances in the Auto Scaling group. Change the number if necessary.

Suggested Answer: D

Community vote distribution

B (96%)

4%

  **princajen**  1 year, 10 months ago

Selected Answer: B

Vote B!

Target.FailedHealthChecks

Verify that the target is listening for traffic on the health check port. You can use the ss command on Linux targets to verify which ports your server is listening on. For Windows targets, you can use the netstat command.

<https://aws.amazon.com/premiumsupport/knowledge-center/elb-fix-failing-health-checks-alb/>

upvoted 14 times

  **[Removed]**  9 months, 3 weeks ago

Selected Answer: B

I would vote for B

upvoted 1 times

  **ronnykapo** 1 year ago

B is correct

upvoted 1 times

  **Olelukoe** 1 year ago

Selected Answer: B

It will fail to start, health check is a next stage



upvoted 1 times

  **caputmundi666** 1 year, 3 months ago

Selected Answer: B

Fix answer to B

upvoted 1 times

  **FKZ** 1 year, 4 months ago

Selected Answer: B

B for sure.

upvoted 1 times

  **braveheart22** 1 year, 4 months ago

B is the way.



upvoted 1 times

  **Idriss10** 1 year, 5 months ago

Selected Answer: D

I vote for D because as mentioned in the question, the ALB and the ASG uses only one EC2 instance.

upvoted 1 times

  **Bhrino** 1 year, 3 months ago

if the maximum number of instances was the issue then the health checks wouldn't fail rather they wouldn't even start so there would be not instances to fail
upvoted 3 times

  **BietTuot** 1 year, 6 months ago

Selected Answer: B

Answer is B

upvoted 1 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: B



bbbbbbb

upvoted 1 times

  **Liongeek** 1 year, 7 months ago

Ans: B



upvoted 2 times

  **[Removed]** 1 year, 8 months ago

Selected Answer: B



It's B.

upvoted 1 times

  **hanguyen89** 1 year, 9 months ago

D em ơi ! hungvv6

upvoted 1 times

  **haxaffee** 1 year, 10 months ago

Selected Answer: B

Vote B. No idea how D is related with failing health checks.

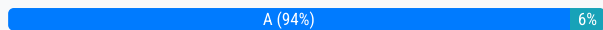
upvoted 4 times

A SysOps administrator has created an AWS Service Catalog portfolio and has shared the portfolio with a second AWS account in the company. The second account is controlled by a different administrator. Which action will the administrator of the second account be able to perform?

- A. Add a product from the imported portfolio to a local portfolio.
- B. Add new products to the imported portfolio.
- C. Change the launch role for the products contained in the imported portfolio.
- D. Customize the products in the imported portfolio.

Suggested Answer: A

Community vote distribution



princajen Highly Voted 2 years, 10 months ago

Selected Answer: A

When you share a portfolio using account-to-account sharing or AWS Organizations, you allow an AWS Service Catalog administrator of another AWS account to import your portfolio into his or her account and distribute the products to end users in that account.

https://docs.aws.amazon.com/servicecatalog/latest/adminguide/catalogs_portfolios_sharing.html

upvoted 15 times

flaacko Most Recent 9 months, 3 weeks ago

Selected Answer: B

You can share portfolios between other AWS accounts, and give the administrators of those accounts the ability to add their own products to your portfolio. This could be useful when you have teams that operate independently, that each deal with creating their own products. You would want their product managers or admins in charge of when new versions are available and to whom they can be provisioned. All you need to do to allow sharing for your portfolio is specify the account id you want to share with (within the service catalog console)

upvoted 1 times

Grodgar 6 months, 2 weeks ago

No, it's not possible.

upvoted 1 times

jipark 1 year, 10 months ago

Selected Answer: A

The administrator of the second account cannot add/customize the imported portfolio created by the first account. second account can only use it.

upvoted 1 times

michaldavid 2 years, 6 months ago

Selected Answer: A

aaaaaaaaa

upvoted 1 times

Liongeek 2 years, 7 months ago

Ans: A

upvoted 1 times

A company has migrated its application to AWS. The company will host the application on Amazon EC2 instances of multiple instance families. During initial testing, a SysOps administrator identifies performance issues on selected EC2 instances. The company has a strict budget allocation policy, so the SysOps administrator must use the right resource types with the performance characteristics to match the workload. What should the SysOps administrator do to meet this requirement?

- A. Purchase regional Reserved Instances (RIs) for immediate cost savings. Review and take action on the EC2 rightsizing recommendations in Cost Explorer. Exchange the RIs for the optimal instance family after rightsizing.
- B. Purchase zonal Reserved Instances (RIs) for the existing instances. Monitor the RI utilization in the AWS Billing and Cost Management console. Make adjustments to instance sizes to optimize utilization.
- C. Review and take action on AWS Compute Optimizer recommendations. Purchase Compute Savings Plans to reduce the cost that is required to run the compute resources.
- D. Review resource utilization metrics in the AWS Cost and Usage Report. Rightsize the EC2 instances. Create On-Demand Capacity Reservations for the rightsized resources.

Suggested Answer: C

Community vote distribution

C (100%)

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


Selected Answer: C

A and B are incorrect as purchasing RI before performing rightsizing is nonsense.

C is the most likely to be correct, but it does not state anything on rightsizing actions. However purchasing Compute Savings Plans allows to make rightsizing later.

D - incorrect, as On-Demand Capacity Reservations has nothing to budget.


upvoted 11 times

 **jipark** 1 year, 4 months ago

I agree :

A, B raise costs, D is only result.

upvoted 1 times

 **kati2k22cz** Highly Voted 2 years, 3 months ago

Selected Answer: C

C


<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recommendations.html>

upvoted 11 times

 **tamng** Most Recent 1 year ago

C is correct

upvoted 1 times

 **FKZ** 1 year, 10 months ago

C

<https://aws.amazon.com/pt/savingsplans/compute-pricing/>



upvoted 1 times

 **MrMLB** 2 years ago

D

The SysOps administrator should review resource utilization metrics in the AWS Cost and Usage Report, rightsize the EC2 instances, and create On-Demand Capacity Reservations for the rightsized resources in order to meet this requirement. This will allow the administrator to identify the EC2 instances that are experiencing performance issues, and resize them to the optimal instance family based on their workload. Creating On-Demand Capacity Reservations will also ensure that the company has access to the necessary compute resources without exceeding its budget allocation. Options A, B, and C are not appropriate for this situation, as they do not involve rightsizing the EC2 instances to match the workload.

upvoted 2 times

  **tamng** 1 year ago

you wrong, C not D

upvoted 1 times

  **Michaldavid** 2 years ago

Selected Answer: C

cccccc

upvoted 2 times

  **Liongeek** 2 years, 1 month ago

Ans: C

upvoted 1 times

A SysOps administrator is tasked with deploying a company's infrastructure as code. The SysOps administrator want to write a single template that can be reused for multiple environments.

How should the SysOps administrator use AWS CloudFormation to create a solution?

- A. Use Amazon EC2 user data in a CloudFormation template.
- B. Use nested stacks to provision resources.
- C. Use parameters in a CloudFormation template.
- D. Use stack policies to provision resources.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **princajen** Highly Voted 👍 1 year, 4 months ago

Selected Answer: C

Reuse templates to replicate stacks in multiple environments

After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production. To make templates reusable, use the parameters, mappings, and conditions sections so that you can customize your stacks when you create them. For example, for your development environments, you can specify a lower-cost instance type compared to your production environment, but all other configurations and settings remain the same.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html#reuse>

upvoted 19 times

🗳️ 👤 **Amonchaiz** Most Recent ⌚ 9 months, 1 week ago

Selected Answer: C

ans: C

upvoted 2 times

🗳️ 👤 **michaldavid** 1 year ago

Selected Answer: C

cccccccc

upvoted 2 times

🗳️ 👤 **Liongeek** 1 year, 1 month ago

Ans: C

upvoted 1 times

🗳️ 👤 **Atown** 1 year, 1 month ago

Selected Answer: C

I believe C

upvoted 2 times

🗳️ 👤 **Milus** 1 year, 2 months ago

It looks like B is correct

upvoted 3 times

🗳️ 👤 **Surferbolt** 1 year, 2 months ago

Selected Answer: C

C. Parameters.

upvoted 2 times

🗳️ 👤 **Gianiluca** 1 year, 3 months ago

D seems correct

upvoted 1 times

🗳️ 👤 **Gianiluca** 1 year, 3 months ago

I meant B as per <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-nested-stacks.html>

upvoted 1 times

  **zolthar_z** 1 year ago

Don't confuse nested stacks with stacksets

upvoted 1 times

A SysOps administrator is responsible for a large fleet of Amazon EC2 instances and must know whether any instances will be affected by upcoming hardware maintenance.

Which option would provide this information with the LEAST administrative overhead?

- A. Deploy a third-party monitoring solution to provide real-time EC2 instance monitoring.
- B. List any instances with failed system status checks using the AWS Management Console.
- C. Monitor AWS CloudTrail for StopInstances API calls.
- D. Review the AWS Personal Health Dashboard.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Hatem08** 7 months ago

Selected Answer: D

dddd D :)

upvoted 1 times

🗳️ 👤 **jipark** 10 months, 3 weeks ago

Selected Answer: D

D : "scheduled change"

"Use the Scheduled changes tab to view upcoming events that might affect your account. These events can include scheduled maintenance activities for services."

upvoted 2 times

🗳️ 👤 **airraid2010** 1 year, 1 month ago

Selected Answer: D

By reviewing the AWS Personal Health Dashboard, the SysOps administrator can quickly and easily identify any instances that will be affected by upcoming hardware maintenance. This information can then be used to take action to mitigate the impact of the maintenance, such as by moving instances to a different Availability Zone.

upvoted 2 times

🗳️ 👤 **piavik** 1 year, 2 months ago

Selected Answer: D

A - possible, but not using AWS services and has admin overhead, so - incorrect

B and C - may show only already failed instances, so - incorrect

D is correct, it shows current and also potential failures on your instances that run on degraded hardware

<https://docs.aws.amazon.com/health/latest/ug/getting-started-health-dashboard.html>

upvoted 4 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

dddddddddd

upvoted 1 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

Ans: D

upvoted 1 times

🗳️ 👤 **princajen** 1 year, 10 months ago

Selected Answer: D

If a hardware malfunction occurs, then Amazon EC2 tags the specific hardware as faulty. Any instances that are running on the hypervisor of the faulty hardware are moved to healthy hardware. During the transition to new hardware, the Amazon EBS-backed instances are stopped and instance store-backed instances are terminated. Amazon EC2 sends a notification through email and to your Personal Health Dashboard informing you of the hardware degradation and of the upcoming instance stop or termination.

<https://aws.amazon.com/premiumsupport/knowledge-center/ec2-linux-degraded-hardware/>
upvoted 4 times

A SysOps administrator is attempting to deploy resources by using an AWS CloudFormation template. An Amazon EC2 instance that is defined in the template fails to launch and produces an `InsufficientInstanceCapacity` error.

Which actions should the SysOps administrator take to resolve this error? (Choose two.)

- A. Create a separate AWS CloudFormation template for the EC2 instance.
- B. Modify the AWS CloudFormation template to not specify an Availability Zone for the EC2 instance.
- C. Modify the AWS CloudFormation template to use a different EC2 instance type.
- D. Use a different Amazon Machine Image (AMI) for the EC2 instance.
- E. Use the AWS CLI's `validate-template` command before creating a stack from the template.

Suggested Answer: BC

Community vote distribution

BC (100%)

  **princajen** Highly Voted 2 years, 4 months ago

Selected Answer: BC

If you're launching an instance, submit a new request without specifying an Availability Zone.

If you're launching an instance, submit a new request using a different instance type (which you can resize at a later stage). For more information, see [Change the instance type](#).

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/troubleshooting-launch.html#troubleshooting-launch-capacity>

upvoted 15 times

  **jipark** 1 year, 4 months ago

"AWS does not currently have enough available On-Demand capacity to fulfill your request."

upvoted 1 times

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

Selected Answer: BC

BC. `InsufficientInstanceCapacity` error message means the current AZ that you wish to deploy your instance in does not have enough capacity, so BC are the best options to circumvent it.

upvoted 7 times

  **tamng** Most Recent 1 year ago


B and C are correct answers

`InsufficientInstanceCapacity` : if you get this error, it means AWS does not have that enough On-Demand capacity in the particular AZ where the instance is launched.

Resolution :

- Wait for few mins before requesting again.
- If requesting more than 1 requests, break down the requests. If you need 5 instances, rather than a single request of 5, request one by one.
- If urgent, submit a request for a different instance type now , which can be resized later.
- Also, can request the EC2 instance in a different AZ

upvoted 1 times

  **piavik** 1 year, 8 months ago

Selected Answer: BC

B and C

upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: BC

B and C

upvoted 1 times

  **Liongeek** 2 years, 1 month ago

Ans: BC

upvoted 1 times

A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses Amazon Route 53 to route traffic.

The company also has a static website that is configured in an Amazon S3 bucket.

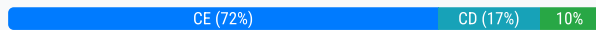
A SysOps administrator must use the static website as a backup to the web application. The failover to the static website must be fully automated.

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

- A. Create a primary failover routing policy record. Configure the value to be the ALB.
- B. Create an AWS Lambda function to switch from the primary website to the secondary website when the health check fails.
- C. Create a primary failover routing policy record. Configure the value to be the ALB. Associate the record with a Route 53 health check.
- D. Create a secondary failover routing policy record. Configure the value to be the static website. Associate the record with a Route 53 health check.
- E. Create a secondary failover routing policy record. Configure the value to be the static website.

Suggested Answer: CD

Community vote distribution



🗳️ 👤 **Nono90** Highly Voted 2 years, 7 months ago

Selected Answer: CE

<https://aws.amazon.com/pt/premiumsupport/knowledge-center/fail-over-s3-r53/>

upvoted 8 times

🗳️ 👤 **FKZ** Highly Voted 2 years, 4 months ago

Selected Answer: CE

No Health Check to secondary fail over.

upvoted 6 times

🗳️ 👤 **Pisces225** Most Recent 1 month, 1 week ago

Selected Answer: AE

Everyone saying CE is wrong, period. If you were not using the ALB then yes, CE would be correct. Pay attention to this which clearly says that if you are using an alias record, as you can with a ELB/ALB, then the health check is configured directly on Route 53. If you were just routing directly to the EC2 instance, then yes, you would need the health check, but that is not the case here.

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

upvoted 1 times

🗳️ 👤 **griggrig** 11 months ago

Selected Answer: CE

No needed for configure health check for S3 static website, only for ALB. Image if you have both health check and both of them unhealthy. In this scenarion no meaning of failover routing.

upvoted 2 times

🗳️ 👤 **mamas_devops** 1 year, 2 months ago

A and E

ALB and S3 are both alias record so no need to create and associate additional health checks. Instead you configuring active-passive failover with one primary and one secondary resource.

Useful links:

<https://repost.aws/questions/QUhIXjnA87QwCWeTjlp1v13g/route-53-alb-alias-record-health-checks#COLFaTiEnLSbiJABZjCX070Q>

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

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

upvoted 2 times

🗳️ 👤 **mamas_devops** 1 year, 2 months ago

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

upvoted 1 times

🗳️ 👤 **Suksay** 1 year, 4 months ago

Selected Answer: CE

For Evaluate target health, turn off the option. When the target is an S3 bucket, the option doesn't wor

upvoted 1 times

🗳️ 👤 **tamng** 1 year, 6 months ago

C E are correct answers

upvoted 1 times

🗳️ 👤 **tamng** 1 year, 6 months ago

Why not CD?

upvoted 1 times

🗳️ 👤 **Bhrino** 2 years, 3 months ago

why would you not associate a health check to the secondary fail over? it doesn't say not to

upvoted 1 times

🗳️ 👤 **Gomer** 2 years, 1 month ago

"for Evaluate target health, turn off the option. When the target is an S3 bucket, the option doesn't work."

upvoted 3 times

🗳️ 👤 **jipark** 1 year, 10 months ago

health check is not clue.

in the chain of Route53 - ALB - EC2 - S3,

Route53 do not need to directly connect to S3. instead, it associate with ALB.

upvoted 1 times

🗳️ 👤 **MrMLB** 2 years, 6 months ago

D E

The combination of actions that will meet these requirements is creating a primary failover routing policy record, configuring the value to be the ALB, and creating a secondary failover routing policy record, configuring the value to be the static website. This allows the SysOps administrator to use the static website as a backup to the web application, and automatically fail over to the static website if the health check fails. The administrator can also associate the record with a Route 53 health check to monitor the health of the web application and trigger the failover if necessary. Option A is not appropriate for this situation, as it does not include the secondary failover routing policy record that is necessary for automatic failover. Option B is not appropriate, as it does not involve configuring the routing policy records, which are necessary for directing traffic to the appropriate website. Option C is not appropriate, as it does not include the secondary failover routing policy record that is necessary for automatic failover.

upvoted 1 times

🗳️ 👤 **noahsark** 2 years, 5 months ago

D is wrong. For Associate with Health Check, choose No.

<https://aws.amazon.com/pt/premiumsupport/knowledge-center/fail-over-s3-r53/>

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: C

C and E

upvoted 3 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: CD

upvoted 2 times

🗳️ 👤 **Capy** 2 years, 7 months ago

C E

read carefully this link. It's mentioned that for the primary routing policy you must specify Health Check YES, instead for the secondary routing failover routing policy with don't specify Health Check

upvoted 2 times

🗳️ 👤 **andreapp** 2 years, 8 months ago

Selected Answer: CD

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html#dns-failover-types-active-passive-one-resource>

upvoted 1 times

🗨️ 👤 **Gomer** 2 years, 1 month ago

"for Evaluate target health, turn off the option. When the target is an S3 bucket, the option doesn't work."

upvoted 1 times

🗨️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: CE

CE. No health check for secondary failover.

upvoted 1 times

🗨️ 👤 **princajen** 2 years, 9 months ago

Selected Answer: CE

Create the failover endpoint

Open the Amazon Route 53 console, and then choose Hosted zones.

Choose the hosted zone that you want to create the record for.

Choose Create record, and input the following:

For Record name, use the same value that you entered for the primary record.

For Record type, choose A – Routes traffic to an IPV4 address and some AWS resources.

For Alias, choose Yes.

Note: Aliases automatically use a time to live (TTL) that matches the alias target.

For Alias Target, choose the S3 bucket that you created previously.

For Routing Policy, choose Failover.

For Failover Record Type, choose Secondary.

For Set ID, enter a name.

Note: The name for the Set ID on your failover endpoint must be different from the name of the Set ID on your primary endpoint.

For Associate with Health Check, choose No.

Choose Create records.

upvoted 2 times

A data analytics application is running on an Amazon EC2 instance. A SysOps administrator must add custom dimensions to the metrics collected by the Amazon CloudWatch agent.

How can the SysOps administrator meet this requirement?

- A. Create a custom shell script to extract the dimensions and collect the metrics using the Amazon CloudWatch agent.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to evaluate the required custom dimensions and send the metrics to Amazon Simple Notification Service (Amazon SNS).
- C. Create an AWS Lambda function to collect the metrics from AWS CloudTrail and send the metrics to an Amazon CloudWatch Logs group.
- D. Create an `append_dimensions` field in the Amazon CloudWatch agent configuration file to collect the metrics.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Hisayuki** Highly Voted 1 year, 4 months ago

Selected Answer: D

Example of CloudWatch agent configuration file:

```
"append_dimensions" : {  
  "ImageId" : "${aws:ImageId}"  
  "InstanceId" : "${aws:InstanceId}"  
}
```

upvoted 16 times

🗳️ 👤 **jipark** 10 months, 3 weeks ago

wow great !!

upvoted 1 times

🗳️ 👤 **princajen** Highly Voted 1 year, 10 months ago

Selected Answer: D

In custom metrics, the `--dimensions` parameter is common. A dimension further clarifies what the metric is and what data it stores. You can have up to 30 dimensions assigned to one metric, and each dimension is defined by a name and value pair

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

upvoted 7 times

🗳️ 👤 **BietTuot** Most Recent 1 year, 6 months ago

Selected Answer: D

Answer is D

upvoted 1 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddddd

upvoted 1 times

A company stores its data in an Amazon S3 bucket. The company is required to classify the data and find any sensitive personal information in its S3 files.

Which solution will meet these requirements?

- A. Create an AWS Config rule to discover sensitive personal information in the S3 files and mark them as noncompliant.
- B. Create an S3 event-driven artificial intelligence/machine learning (AI/ML) pipeline to classify sensitive personal information by using Amazon Rekognition.
- C. Enable Amazon GuardDuty. Configure S3 protection to monitor all data inside Amazon S3.
- D. Enable Amazon Macie. Create a discovery job that uses the managed data identifier.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **princajen** Highly Voted 👍 1 year, 10 months ago

Selected Answer: D

To discover sensitive data with Amazon Macie, you create and run sensitive data discovery jobs. A sensitive data discovery job analyzes objects in Amazon Simple Storage Service (Amazon S3) buckets to determine whether the objects contain sensitive data, and it provides detailed reports of the sensitive data that it finds and the analysis that it performs. By creating and running jobs, you can automate discovery, logging, and reporting of sensitive data in S3 buckets.

<https://docs.aws.amazon.com/macie/latest/user/data-classification.html>

upvoted 8 times

🗳️ 👤 **[Removed]** Most Recent ⌚ 11 months, 3 weeks ago

D - 100%

upvoted 2 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddddd

upvoted 2 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

Ans: D

upvoted 2 times

🗳️ 👤 **Surferbolt** 1 year, 8 months ago

Selected Answer: D

D. Amazon Macie searches S3 buckets for sensitive personal information.

upvoted 1 times

A company hosts a web portal on Amazon EC2 instances. The web portal uses an Elastic Load Balancer (ELB) and Amazon Route 53 for its public DNS service.

The ELB and the EC2 instances are deployed by way of a single AWS CloudFormation stack in the us-east-1 Region. The web portal must be highly available across multiple Regions.

Which configuration will meet these requirements?

- A. Deploy a copy of the stack in the us-west-2 Region. Create a single start of authority (SOA) record in Route 53 that includes the IP address from each ELB. Configure the SOA record with health checks. Use the ELB in us-east-1 as the primary record and the ELB in us-west-2 as the secondary record.
- B. Deploy a copy of the stack in the us-west-2 Region. Create an additional A record in Route 53 that includes the ELB in us-west-2 as an alias target. Configure the A records with a failover routing policy and health checks. Use the ELB in us-east-1 as the primary record and the ELB in us-west-2 as the secondary record.
- C. Deploy a new group of EC2 instances in the us-west-2 Region. Associate the new EC2 instances with the existing ELB, and configure load balancer health checks on all EC2 instances. Configure the ELB to update Route 53 when EC2 instances in us-west-2 fail health checks.
- D. Deploy a new group of EC2 instances in the us-west-2 Region. Configure EC2 health checks on all EC2 instances in each Region. Configure a peering connection between the VPCs. Use the VPC in us-east-1 as the primary record and the VPC in us-west-2 as the secondary record.

Suggested Answer: A

Community vote distribution

B (100%)

  **Balliache520505** Highly Voted 1 year, 9 months ago

The answer is B. When you create a hosted zone, Route 53 automatically creates a name server (NS) record and a start of authority (SOA) record for the zone. <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/migrate-dns-domain-in-use.html#migrate-dns-create-hosted-zone>
upvoted 12 times

  **Benly** Most Recent 8 months ago

Selected Answer: B

Why is the answer A?
upvoted 1 times

  **[Removed]** 9 months, 3 weeks ago

Selected Answer: B

B since SOA record is not appropriate for this
upvoted 1 times

  **jipark** 10 months, 3 weeks ago

Selected Answer: B

deploy "copy stack", then use "Alias"
upvoted 2 times

  **sisover** 11 months ago

What happens when people vote for B and the question shows different answer as correct one ?
upvoted 1 times

  **Andrew_A** 1 year ago

Selected Answer: B

Option A involves a start of authority (SOA) record, which is not appropriate for failover configuration.
upvoted 4 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: B

bbbbbbbbb
upvoted 2 times

  **Liongeek** 1 year, 7 months ago

Ans: B

upvoted 2 times

🗨️ 👤 **Shruti09753** 1 year, 8 months ago

Selected Answer: B

B should be the answer

upvoted 2 times

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

Selected Answer: B

B is the answer

upvoted 1 times

🗨️ 👤 **hippius** 1 year, 8 months ago

Selected Answer: B

Sure It's B

upvoted 2 times

🗨️ 👤 **Gianiluca** 1 year, 9 months ago

Selected Answer: B

Agree it's B

upvoted 1 times

🗨️ 👤 **AAAaat** 1 year, 9 months ago

Selected Answer: B

B. the doc also state you rarely would ever need to change anything about SOA which is automoatically created.

upvoted 2 times

🗨️ 👤 **princajen** 1 year, 9 months ago

Selected Answer: B

No matter how I read it, I always go for B!

upvoted 1 times

🗨️ 👤 **andrerkn** 1 year, 9 months ago

Selected Answer: B

Going with B

upvoted 1 times

🗨️ 👤 **haxaffee** 1 year, 9 months ago

Selected Answer: B

It's not possible to create or add failover to SOA records and you can also not Alias target them. Wrong in so many ways. Please read https://en.wikipedia.org/wiki/SOA_record

B.

upvoted 1 times

🗨️ 👤 **kati2k22cz** 1 year, 9 months ago

Selected Answer: B

Guess is B Because the use of routing policy and health checks

upvoted 2 times

A SysOps administrator is investigating why a user has been unable to use RDP to connect over the internet from their home computer to a bastion server running on an Amazon EC2 Windows instance.

Which of the following are possible causes of this issue? (Choose two.)

- A. A network ACL associated with the bastion's subnet is blocking the network traffic.
- B. The instance does not have a private IP address.
- C. The route table associated with the bastion's subnet does not have a route to the internet gateway.
- D. The security group for the instance does not have an inbound rule on port 22.
- E. The security group for the instance does not have an outbound rule on port 3389.

Suggested Answer: AC

Community vote distribution


AC (100%)

 **haxaffee** Highly Voted 2 years, 3 months ago

Selected Answer: AC


Correct.

upvoted 8 times

 **Geek20** Most Recent 8 months, 1 week ago

A&C are correct

upvoted 1 times

 **jipark** 1 year, 4 months ago

Selected Answer: AC

not D : RDP uses port 3389

upvoted 2 times

 **fazlur21** 1 year, 6 months ago

ac correct

upvoted 1 times

 **MrMLB** 2 years ago

The possible causes of this issue are:

- A. A network ACL associated with the bastion's subnet is blocking the network traffic.
- C. The route table associated with the bastion's subnet does not have a route to the internet gateway.
- E. The security group for the instance does not have an outbound rule on port 3389.

upvoted 4 times

 **student2020** 1 year, 11 months ago

E is wrong, the rule for 3389 must be inbound NOT outbound

upvoted 3 times

 **AgboolaKun** 7 months, 2 weeks ago

E is wrong because security groups are stateful. You only need to specify an inbound rule, outbound allows all traffic.

upvoted 1 times

 **Jamshif01** 2 years, 1 month ago

A - can be

B - he is connecting via public anyway this is irrelevant

C - yes it should have access to internet so he can connect outside

D - this is for linux only

E - can be

so it's either AE or CE

I vote for CE

upvoted 2 times

  **Jamshif01** 2 years, 1 month ago

I was wrong about E it says outbound
so the answer is AC 100%

upvoted 1 times

  **Liongeek** 2 years, 1 month ago

Selected Answer: AC

Ans: A&C

upvoted 2 times

  **Gianiluca** 2 years, 3 months ago

C and E - <https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/troubleshoot-connect-windows-instance.html>

upvoted 2 times

  **Liongeek** 2 years, 1 month ago



Outbound is not a concern here.

upvoted 1 times

  **Surferbolt** 2 years, 2 months ago

Not E. Security Groups are stateful, inbound rules will apply to outbound rules.

upvoted 2 times

  **nikkoe** 2 years, 3 months ago

AC, outbound traffic has nothing to do with this

upvoted 2 times

A SysOps administrator is examining the following AWS CloudFormation template:

```
AWSTemplateFormatVersion: '2010-09-09'
Description: 'Creates an EC2 Instance'
Resources:
  EC2Instance:
    Type: AWS::EC2::Instance
    Properties:
      ImageId: ami-79fd7eee
      InstanceType: m5n.large
      SubnetId: subnet-1abc3d3fg
      PrivateDnsName: ip-10-24-34-0.ec2.internal
      Tags:
        - Key: Name
          Value: !Sub "${AWS::StackName} Instance"
```

Why will the stack creation fail?

- A. The Outputs section of the CloudFormation template was omitted.
- B. The Parameters section of the CloudFormation template was omitted.
- C. The PrivateDnsName cannot be set from a CloudFormation template.
- D. The VPC was not specified in the CloudFormation template.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Liongeek** Highly Voted 1 year, 7 months ago

Ans: C

PrivateDNSEndpoint is an output, not a parameter

upvoted 9 times

🗨️ 👤 **haxaffee** Highly Voted 1 year, 9 months ago

Selected Answer: C

C is correct. See <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-ec2-instance.html> Only available is PrivateDnsNameOptions.

upvoted 7 times

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

"PrivateDnsNameOptions": PrivateDnsNameOptions,

upvoted 2 times

🗨️ 👤 **wendy971** Most Recent 4 months, 2 weeks ago

Selected Answer: C

AWS CloudFormation templates only have Resource as a required field

upvoted 1 times

A new application runs on Amazon EC2 instances and accesses data in an Amazon RDS database instance. When fully deployed in production, the application fails. The database can be queried from a console on a bastion host. When looking at the web server logs, the following error is repeated multiple times:

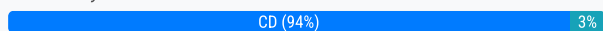
*** Error Establishing a Database Connection

Which of the following may be causes of the connectivity problems? (Choose two.)

- A. The security group for the database does not have the appropriate egress rule from the database to the web server.
- B. The certificate used by the web server is not trusted by the RDS instance.
- C. The security group for the database does not have the appropriate ingress rule from the web server to the database.
- D. The port used by the application developer does not match the port specified in the RDS configuration.
- E. The database is still being created and is not available for connectivity.

Suggested Answer: CE

Community vote distribution



🗳️ **sxti** Highly Voted 2 years ago

Selected Answer: CD

CD because you are not stupid
upvoted 12 times

🗳️ **Vinsmoke** 1 year, 8 months ago

unfortunately, i decided to ignore that a bastion had connectivity.
upvoted 2 times

🗳️ **Aamee** Most Recent 8 months ago

Selected Answer: CD

C & D straightfoward..
upvoted 1 times

🗳️ **auxwww** 1 year ago

Selected Answer: CD

Bastion host is connecting, so DB is up! Which leaves C and D
upvoted 3 times

🗳️ **komorebi** 1 year, 1 month ago

Selected Answer: CD

CD for sure
upvoted 1 times

🗳️ **Nowon** 1 year, 3 months ago

Selected Answer: CD

Because of bastion host can queries to database
upvoted 1 times

🗳️ **Suksay** 1 year, 4 months ago

Selected Answer: CD

Not E because in subject : The database can be queried from a console on a bastion host.
upvoted 1 times

🗳️ **Passexam4sure_com** 1 year, 8 months ago

Selected Answer: CD

C. The security group for the database does not have the appropriate ingress rule from the web server to the database.
D. The port used by the application developer does not match the port specified in the RDS configuration.
upvoted 3 times

🗳️ **DennisRichard** 1 year, 8 months ago

Who behind the scenes is marking these 'correct' options? The question clearly states that the bastion host is able to query the database and yet the 'correct' answer is 'instance is still being created'?!?!?! Why Lord Why? hahaha

This question bank is making me more nervous about the exam than I should be feeling.

upvoted 2 times

🗳️ 👤 **sidneyalpagel** 1 year, 8 months ago

CD because the application is fully deployed. So... database not ready doesnt make sense to me

upvoted 1 times

🗳️ 👤 **AShahine21** 1 year, 12 months ago

C and E

For D, I don't know who is stupid enough to use the wrong Database RDS Port.

upvoted 2 times

🗳️ 👤 **xSohox** 1 year, 10 months ago

You need to read questions more attentively.

"The database can be queried from a console on a bastion host."

upvoted 3 times

🗳️ 👤 **wh1t4k3r** 1 year, 9 months ago

For these tests I strongly advise you to base your answers on technicality and not personal experience.

upvoted 1 times

🗳️ 👤 **wh1t4k3r** 1 year, 9 months ago

E is not valid since the access is possible through the bastion.

upvoted 1 times

🗳️ 👤 **piavik** 2 years, 2 months ago

Selected Answer: CD

A and D

A - incorrect, database does not connect to web application

B - incorrect, web server cert is not relevant to database connection

E - incorrect, all is working from bastion host, so database is OK

upvoted 2 times

🗳️ 👤 **squeeze_talus0y** 2 years, 5 months ago

Selected Answer: CD

C and D

upvoted 4 times

🗳️ 👤 **MrMLB** 2 years, 6 months ago

A and C are possible causes of the connectivity problems.

A is a possible cause because the security group for the database controls which other resources can access the database. If the security group does not have the appropriate egress rule from the database to the web server, the web server will not be able to access the database.

C is a possible cause because the security group for the database also controls which other resources can access the database. If the security group does not have the appropriate ingress rule from the web server to the database, the web server will not be able to access the database.

B, D, and E are not likely causes of the connectivity problems. The certificate used by the web server is not relevant to the connectivity between the web server and the RDS instance. The port used by the application developer does not need to match the port specified in the RDS configuration, and the database being created is not relevant to the connectivity problems.

upvoted 1 times

🗳️ 👤 **DennisRichard** 1 year, 8 months ago

There are no 'egress' rules in a security group. You must be thinking of NACLs instead.

upvoted 2 times

🗳️ 👤 **Bhrino** 2 years, 3 months ago

is c and D because it indicates that there is an inbound connection issue which rules out A and what you said about d with the port not needing to match is just wrong. This can cause some inbound connectivity issues.

upvoted 1 times

🗳️ 👤 **marcoeu** 1 year, 9 months ago

A is not possible for me because security group rules are statefull. If inbound rules are allowed then outbound rules in the same ports are allowed.

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: C

C and D

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Selected Answer: CD

Ans: C&D

upvoted 1 times

🗨️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: CD

CD. Database can be queried from Bastion, so E is out. Security groups are stateful, so you don't have to bother with the egress rules in this situation, as long as you have the proper ingress rule.

upvoted 3 times

🗨️ 👤 **hippius** 2 years, 8 months ago

Selected Answer: CD

CD because the database can be queried from the bastion

upvoted 1 times

A compliance team requires all administrator passwords for Amazon RDS DB instances to be changed at least annually.

Which solution meets this requirement in the MOST operationally efficient manner?

- A. Store the database credentials in AWS Secrets Manager. Configure automatic rotation for the secret every 365 days.
- B. Store the database credentials as a parameter in the RDS parameter group. Create a database trigger to rotate the password every 365 days.
- C. Store the database credentials in a private Amazon S3 bucket. Schedule an AWS Lambda function to generate a new set of credentials every 365 days.
- D. Store the database credentials in AWS Systems Manager Parameter Store as a secure string parameter. Configure automatic rotation for the parameter every 365 days.

Suggested Answer: A

Community vote distribution

A (100%)

  **FKZ** Highly Voted 2 years, 3 months ago

Selected Answer: 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. For more information, see [What is AWS Secrets Manager?](#) in the AWS Secrets Manager User Guide.



upvoted 7 times

  **Aamee** Most Recent 8 months ago

Selected Answer: A

It should be only A since it's specifically asked for 'Operationally efficient', means a solution that can contain with as less no. of services as possible and can deliver efficiency too.

upvoted 1 times



  **noahsark** 2 years, 5 months ago

Selected Answer: A

Rotate Amazon RDS database credentials automatically with AWS Secrets Manager



<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>

upvoted 3 times

  **jipark** 1 year, 10 months ago

same questions from Dev Ops

upvoted 1 times

  **MrMLB** 2 years, 6 months ago

D

The most operationally efficient solution would be to store the database credentials in AWS Systems Manager Parameter Store as a secure string parameter, and then configure automatic rotation for the parameter every 365 days. This way, the credentials will be securely stored and automatically rotated as required by the compliance team. Options A and C both involve storing the credentials in different locations and using different methods for rotating the credentials, which would not be as operationally efficient. Option B involves using a database trigger to rotate the password, but this would require additional setup and maintenance, and may not be as reliable as using automatic rotation in AWS Systems Manager.

upvoted 1 times

  **noahsark** 2 years, 5 months ago

A. Store the database credentials in AWS Secrets Manager. Configure automatic rotation for the secret every 365 days.

<https://aws.amazon.com/blogs/security/rotate-amazon-rds-database-credentials-automatically-with-aws-secrets-manager/>

upvoted 2 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaaaaa

upvoted 2 times

  **Liongeek** 2 years, 7 months ago

Ams: A

upvoted 1 times

  **kati2k22cz** 2 years, 9 months ago

Selected Answer: A

A is correct. We can confirm here

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

upvoted 4 times

A SysOps administrator is responsible for managing a fleet of Amazon EC2 instances. These EC2 instances upload build artifacts to a third-party service. The third-party service recently implemented a strict IP allow list that requires all build uploads to come from a single IP address. What change should the systems administrator make to the existing build fleet to comply with this new requirement?

- A. Move all of the EC2 instances behind a NAT gateway and provide the gateway IP address to the service.
- B. Move all of the EC2 instances behind an internet gateway and provide the gateway IP address to the service.
- C. Move all of the EC2 instances into a single Availability Zone and provide the Availability Zone IP address to the service.
- D. Move all of the EC2 instances to a peered VPC and provide the VPC IP address to the service.

Suggested Answer: A

Community vote distribution

A (100%)

🗲️ 👤 **kati2k22cz** Highly Voted 👍 1 year, 3 months ago

Selected Answer: A

NAT Gateway is the choice. Letter A

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

upvoted 10 times

🗲️ 👤 **Bhrino** Most Recent 🔍 9 months, 3 weeks ago

Selected Answer: A

Option A is the only option that would provide a single ip address

upvoted 4 times

🗲️ 👤 **michaldavid** 1 year ago

Selected Answer: A

aaaaaaaaa

upvoted 1 times

🗲️ 👤 **Liongeek** 1 year, 1 month ago

Ans: A

upvoted 1 times

A company uses an Amazon CloudFront distribution to deliver its website. Traffic logs for the website must be centrally stored, and all data must be encrypted at rest.

Which solution will meet these requirements?

- A. Create an Amazon OpenSearch Service (Amazon Elasticsearch Service) domain with internet access and server-side encryption that uses the default AWS managed customer master key (CMK). Configure CloudFront to use the Amazon OpenSearch Service (Amazon Elasticsearch Service) domain as a log destination.
- B. Create an Amazon OpenSearch Service (Amazon Elasticsearch Service) domain with VPC access and server-side encryption that uses AES-256. Configure CloudFront to use the Amazon OpenSearch Service (Amazon Elasticsearch Service) domain as a log destination.
- C. Create an Amazon S3 bucket that is configured with default server-side encryption that uses AES-256. Configure CloudFront to use the S3 bucket as a log destination.
- D. Create an Amazon S3 bucket that is configured with no default encryption. Enable encryption in the CloudFront distribution, and use the S3 bucket as a log destination.

Suggested Answer: C

Community vote distribution

C (100%)

  **Balliache520505** Highly Voted 2 years, 9 months ago

The answer is C. <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/AccessLogs.html>
upvoted 7 times

  **jipark** 1 year, 10 months ago

S3 contains Encryption of Data without additional KMS solutions.
upvoted 1 times

  **gehadg** Most Recent 8 months ago

Option C: Storing CloudFront logs in an Amazon S3 bucket with server-side encryption using AES-256 meets the requirement for central log storage and data encryption at rest. Amazon S3 provides server-side encryption, and configuring CloudFront to log directly to this bucket is a common and effective way to handle CloudFront logs securely.
upvoted 1 times

  **BietTuot** 2 years, 6 months ago

Selected Answer: C

Answer is C
upvoted 3 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: C

ccccccc
upvoted 1 times

  **Surferbolt** 2 years, 8 months ago

Selected Answer: C

C. CloudFront logs can be sent to an S3 bucket, and S3 buckets can be encrypted.
upvoted 2 times

  **kati2k22cz** 2 years, 9 months ago

Selected Answer: C

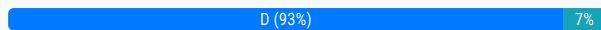
C.
here some references
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html>
<https://stackoverflow.com/questions/52560188/are-my-s3-objects-encrypted-at-rest-or-not>
upvoted 2 times

An organization created an Amazon Elastic File System (Amazon EFS) volume with a file system ID of fs-85ba41fc, and it is actively used by 10 Amazon EC2 hosts. The organization has become concerned that the file system is not encrypted. How can this be resolved?

- A. Enable encryption on each host's connection to the Amazon EFS volume. Each connection must be recreated for encryption to take effect.
- B. Enable encryption on the existing EFS volume by using the AWS Command Line Interface.
- C. Enable encryption on each host's local drive. Restart each host to encrypt the drive.
- D. Enable encryption on a newly created volume and copy all data from the original volume. Reconnect each host to the new volume.

Suggested Answer: D

Community vote distribution



jkwek Highly Voted 3 years, 9 months ago

D is correct.

<https://docs.aws.amazon.com/efs/latest/ug/encryption.html>

Amazon EFS supports two forms of encryption for file systems, encryption of data in transit and encryption at rest. You can enable encryption of data at rest when creating an Amazon EFS file system. You can enable encryption of data in transit when you mount the file system.

upvoted 19 times

rcptryk Most Recent 7 months, 2 weeks ago

Selected Answer: D

Encryption at rest can't be modified.

<https://docs.aws.amazon.com/efs/latest/ug/encryption-at-rest.html#:~:text=You%20can%20create,encrypted%20file%20system.>

upvoted 1 times

gehadg 8 months ago

Encryption at Rest: Amazon EFS encryption at rest can only be enabled at the time of creation. It is not possible to enable encryption on an existing EFS file system after it has been created. Therefore, to have an encrypted EFS volume, you must create a new EFS volume with encryption enabled and then migrate the data from the old volume.

Other Options:

Option A: Encryption on each host's connection to EFS does not exist as an option. Encryption in transit is handled by EFS using TLS automatically.

Option B: There is no command to enable encryption on an existing EFS volume.

Option C: Encrypting each host's local drive does not impact the EFS volume's encryption state.

Correct Answer: D

upvoted 2 times

james2033 1 year, 4 months ago

Selected Answer: D

A --> encrypt connection --> false.

B --> encrypt exist EFS volume (need create new, then migrate to new one) --> false

C --> Encrypt each's host local drive --> No, need create new one, then encrypt new one, then migrate --> causes C wrong.

D --> It is admitted practice with encryption: Create new one, encrypt new one, migrate to that.

upvoted 2 times

noahsark 2 years, 5 months ago

Selected Answer: D

Enable encryption on a newly created volume and copy all data from the original volume. Reconnect each host to the new volume.

<https://aws.amazon.com/premiumsupport/knowledge-center/efs-turn-on-encryption-at-rest/>



upvoted 2 times

  **squeeze_talus0y** 2 years, 5 months ago

Selected Answer: D

Encryption at rest can't be enabled after the EFS volume has been created.

upvoted 3 times

  **MrMLB** 2 years, 6 months ago

Selected Answer: B

B

The best solution for resolving this issue is to enable encryption on the existing EFS volume by using the AWS Command Line Interface. This will allow the organization to encrypt the data on the file system without having to recreate any connections or copy any data. Option A involves enabling encryption on each host's connection to the EFS volume, but this would require each connection to be recreated for encryption to take effect. Option C involves enabling encryption on each host's local drive, but this would not encrypt the data on the EFS volume. Option D involves creating a new volume and copying all data from the original volume, but this would be time-consuming and would require each host to be reconnected to the new volume. Enabling encryption on the existing EFS volume using the AWS CLI is the most efficient and effective solution.

upvoted 1 times

  **noahsark** 2 years, 5 months ago

B is wrong. D is correct.

Enable encryption on a newly created volume and copy all data from the original volume. Reconnect each host to the new volume.

<https://aws.amazon.com/premiumsupport/knowledge-center/efs-turn-on-encryption-at-rest/>

upvoted 2 times



  **foreverlearner** 2 years, 6 months ago

I noticed a few (wrong) answers like this, and they looked like if they were generated by ChatGPT. So I tried it, and it basically gave me the same answer (just slightly longer but most words are exactly the same).

Kind of surprised that ChatGPT is wrong, though, as the AWS doc clearly says "Once you create an EFS file system, you cannot change its encryption setting. This means that you cannot modify an unencrypted file system to make it encrypted. Instead, you need to create a new, encrypted file system." (<https://docs.aws.amazon.com/efs/latest/ug/encryption-at-rest.html>)

Kind of fun that, if you tell ChatGPT the right answer, it apologizes, it admits it's wrong, and also provides with instructions on how to copy the files :D

upvoted 2 times

  **MrMLB** 2 years, 6 months ago

B

The best solution for resolving this issue is to enable encryption on the existing EFS volume by using the AWS Command Line Interface. This will allow the organization to encrypt the data on the file system without having to recreate any connections or copy any data. Option A involves enabling encryption on each host's connection to the EFS volume, but this would require each connection to be recreated for encryption to take effect. Option C involves enabling encryption on each host's local drive, but this would not encrypt the data on the EFS volume. Option D involves creating a new volume and copying all data from the original volume, but this would be time-consuming and would require each host to be reconnected to the new volume. Enabling encryption on the existing EFS volume using the AWS CLI is the most efficient and effective solution.

upvoted 1 times

  **noahsark** 2 years, 5 months ago

B is wrong. D is correct.

Enable encryption on a newly created volume and copy all data from the original volume. Reconnect each host to the new volume.

<https://aws.amazon.com/premiumsupport/knowledge-center/efs-turn-on-encryption-at-rest/>

upvoted 1 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddddd

upvoted 1 times

  **Masoud11** 2 years, 7 months ago

Selected Answer: D

100% D

upvoted 1 times

🗳️ 👤 **Starboy** 2 years, 9 months ago

D is correct as you can't encrypt after the creation of EFS volume.

upvoted 1 times

🗳️ 👤 **Finger41** 3 years, 1 month ago

Selected Answer: D

D -<https://docs.aws.amazon.com/efs/latest/ug/encryption-at-rest.html>

upvoted 2 times

🗳️ 👤 **szl0144** 3 years, 5 months ago

vote D

upvoted 1 times

🗳️ 👤 **ngthien041292** 3 years, 7 months ago

Selected Answer: D

Vote D

upvoted 3 times

🗳️ 👤 **jkwek** 3 years, 8 months ago

The reasoning here for answer D is there is no details for existing data encryption. So to play safe, better encrypt then copy data over.

upvoted 3 times

🗳️ 👤 **RicardoD** 3 years, 9 months ago

D is the answer

upvoted 2 times

A company uses an AWS Service Catalog portfolio to create and manage resources. A SysOps administrator must create a replica of the company's existing AWS infrastructure in a new AWS account.

What is the MOST operationally efficient way to meet this requirement?

- A. Create an AWS CloudFormation template to use the AWS Service Catalog portfolio in the new AWS account.
- B. In the new AWS account, manually create an AWS Service Catalog portfolio that duplicates the original portfolio.
- C. Run an AWS Lambda function to create a new AWS Service Catalog portfolio based on the output of the DescribePortfolio API operation.
- D. Share the AWS Service Catalog portfolio with the new AWS account. Import the portfolio into the new AWS account.

Suggested Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **kati2k22cz** Highly Voted 👍 1 year, 3 months ago

Selected Answer: D

D

https://docs.aws.amazon.com/servicecatalog/latest/adminguide/catalogs_portfolios_sharing.html

upvoted 9 times

🗲️ 👤 **MrMLB** Most Recent 🕒 1 year ago

Selected Answer: D

D it is

upvoted 3 times

🗲️ 👤 **michaldavid** 1 year ago

Selected Answer: D

ddddddd

upvoted 1 times

A SysOps administrator must manage the security of an AWS account. Recently, an IAM user's access key was mistakenly uploaded to a public code repository.

The SysOps administrator must identify anything that was changed by using this access key.

How should the SysOps administrator meet these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to send all IAM events to an AWS Lambda function for analysis.
- B. Query Amazon EC2 logs by using Amazon CloudWatch Logs Insights for all events initiated with the compromised access key within the suspected timeframe.
- C. Search AWS CloudTrail event history for all events initiated with the compromised access key within the suspected timeframe.
- D. Search VPC Flow Logs for all events initiated with the compromised access key within the suspected timeframe.

Suggested Answer: C

Community vote distribution

C (100%)

 **kati2k22cz** Highly Voted 2 years, 9 months ago

Selected Answer: C

C

"You can troubleshoot operational and security incidents over the past 90 days in the CloudTrail console by viewing Event history."

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/view-cloudtrail-events.html>

upvoted 10 times

 **gehadg** Most Recent 8 months ago

AWS CloudTrail records all API activity in an AWS account, including actions taken by IAM users and their access keys. By using CloudTrail, the SysOps administrator can track exactly what actions were performed with the compromised access key, including changes to resources or account configurations, and view a detailed log of events within the specified timeframe.

Other Options:

Option A: Creating an EventBridge (formerly CloudWatch Events) rule for IAM events would not provide historical data about the actions already taken with the compromised key. This setup would only capture future IAM events.

Option B: CloudWatch Logs Insights does not provide logs of all actions taken by IAM users and access keys across services. It is mainly used for querying logs stored in CloudWatch, such as application logs.

Option D: VPC Flow Logs track network traffic but do not log specific API actions, making them unsuitable for identifying API events or actions taken with an IAM access key.

Correct Answer: C

upvoted 3 times

 **jipark** 1 year, 10 months ago

Selected Answer: C

"all event" is key word for answer.

upvoted 1 times

 **jipark** 1 year, 10 months ago

sorry for mistype - for all event of CloudTrail (not EC2 log)

upvoted 1 times

 **BietTuot** 2 years, 6 months ago

Selected Answer: C

answer is C.

upvoted 1 times

 **michaldavid** 2 years, 6 months ago

Selected Answer: C

cccccc

upvoted 1 times

A company runs a retail website on multiple Amazon EC2 instances behind an Application Load Balancer (ALB). The company must secure traffic to the website over an HTTPS connection.

Which combination of actions should a SysOps administrator take to meet these requirements? (Choose two.)

- A. Attach the certificate to each EC2 instance.
- B. Attach the certificate to the ALB.
- C. Create a private certificate in AWS Certificate Manager (ACM).
- D. Create a public certificate in AWS Certificate Manager (ACM).
- E. Export the certificate, and attach it to the website.

Suggested Answer: BD

Community vote distribution

BD (100%)

🗳️ 👤 **jipark** 10 months, 3 weeks ago

Selected Answer: BD

my company do this way :

- create public certificate (call from public internet)
- install certificate on L4 (ALB)

upvoted 4 times

🗳️ 👤 **arjundevops** 1 year ago

Selected Answer: BD

B and D is Correct

upvoted 2 times

🗳️ 👤 **Nrn143** 1 year ago

B and DR correct

upvoted 3 times

🗳️ 👤 **braveheart22** 1 year, 4 months ago

B&D very correct.

upvoted 3 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: BD

B and D

upvoted 1 times

🗳️ 👤 **kati2k22cz** 1 year, 9 months ago

Selected Answer: BD

BD

show article about certificates

<https://www.sslls.com/blog/the-difference-between-aws-certificates-and-public-ca-certificates/>

upvoted 3 times

SIMULATION -

Instructions -

If the copy-paste functionality is not working in your environment, refer to the instructions file on the VM desktop and use Ctrl+C, Ctrl+V or Command-C, Command-V.

Configure Amazon EventBridge to meet the following requirements.

1. Use the us-east-2 Region for all resources.
2. Unless specified below, use the default configuration settings.
3. Use your own resource naming unless a resource name is specified below.
4. Ensure all Amazon EC2 events in the default event bus are replayable for the past 45 days.
5. Create a rule named RunFunction to send the exact message {"name":"example"} every 15 minutes to an existing AWS Lambda function named LogEventFunction
6. Create a rule named SpotWarning to send a notification to a new standard Amazon SNS topic named TopicEvents whenever an Amazon EC2 Spot Instance is interrupted. Do NOT create any topic subscriptions. The notification must match the following structure:

Input path:

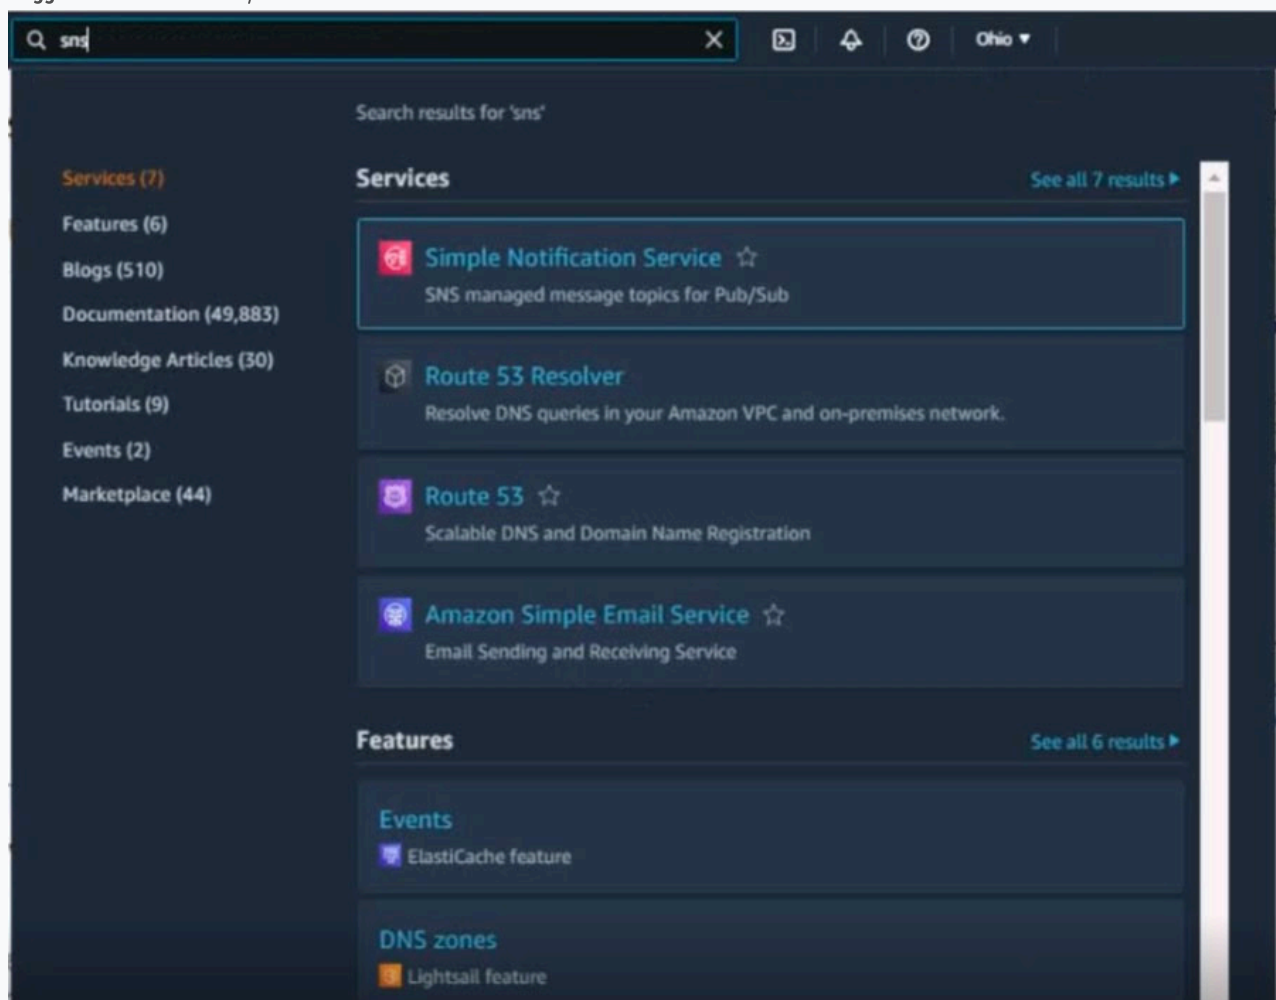
```
{`instance`:`detail.instance-id`}
```

Input template:

```
`The EC2 Spot Instance <instance> has been interrupted.`
```

Important: Click the Next button to complete this lab and continue to the next lab. Once you click the Next button, you will NOT be able to return to this lab.

Suggested Answer: See explanation below.



Amazon Simple Notification Service

Pub/sub messaging for microservices and serverless applications.

Amazon SNS is a highly available, durable, secure, fully managed pub/sub messaging service that enables you to decouple microservices, distributed systems, and event-driven serverless applications. Amazon SNS provides topics for high-throughput, push-based, many-to-many messaging.

Create topic

Topic name

A topic is a message channel. When you publish a message to a topic, it fans out the message to all subscribed endpoints.

[Next step](#)[Start with an overview](#)[Amazon SNS](#) > [Topics](#) > [Create topic](#)

Create topic

Details

Type [Info](#)

Topic type cannot be modified after topic is created

☐ FIFO (first-in, first-out)

- Strictly-preserved message ordering
- Exactly-once message delivery
- High throughput, up to 300 publishes/second
- Subscription protocols: SQS

☒ Standard

- Best-effort message ordering
- At-least once message delivery
- Highest throughput in publishes/second
- Subscription protocols: SQS, Lambda, HTTP, SMS, email, mobile application endpoints

Name

Maximum 256 characters. Can include alphanumeric characters, hyphens (-) and underscores (_).

Display name - optional

To use this topic with SMS subscriptions, enter a display name. Only the first 10 characters are displayed in an SMS message. [Info](#)

Maximum 100 characters.

► Encryption - optional

Amazon SNS provides in-transit encryption by default. Enabling server-side encryption adds at-rest encryption to your topic.

► Access policy - optional

This policy defines who can access your topic. By default, only the topic owner can publish or subscribe to the topic. [Info](#)

► Delivery retry policy (HTTP/S) - optional

The policy defines how Amazon SNS retries failed deliveries to HTTP/S endpoints. To modify the default settings, expand this section. [Info](#)

► Tags - optional

A tag is a metadata label that you can assign to an Amazon SNS topic. Each tag consists of a key and an optional value. You can use tags to search and filter your topics and track your costs. [Learn more](#)

[Cancel](#)[Create topic](#)

Dashboard

Topics

Subscriptions

▼ Mobile

Text messaging (SMS)

Origination numbers

Amazon SNS > Topics > TopicEvents

TopicEvents

Edit

Delete

Publish message

Details

Name

TopicEvents

Display name

-

ARN

arn:aws:sns:us-east-2:082772415256:TopicEvents

Topic owner

082772415256

Type

Standard

Subscriptions

Access policy

Delivery retry policy (HTTP/S)

Delivery status logging

Encryption

Tags

Subscriptions (0)

Edit

Delete

Request confirmation

Confirm subscription

Create subscription

Q Search

< 1 >

⊙

ID

▼

Endpoint

▼

Status

▼

Protocol

▲

Q eventbridge

×

🔍

🔔

🔗

Ohio ▼

Search results for 'eventbridge'

Services (1)

Features (1)

Blogs (87)

Documentation (11,444)

Knowledge Articles (30)

Marketplace (11)

Services



Amazon EventBridge ☆

Serverless event bus that connects application data from your own apps, SaaS, and ...

Features

Schemas



Amazon EventBridge feature

Amazon
EventBridge

×

🔔 Let us know what you think!

We recently made changes to our console experience (sandbox and rule creation flow) and we'd love to hear what you think!

Provide Feedback

×

▼ Getting started

Learn

Sandbox **New**

▼ Events

Event buses

Rules

Global endpoints **New**

Archives

Replays

▼ Integration

Partner event sources

API destinations

▼ Schema registry

Schemas

Application Integration

Amazon EventBridge

Build event-driven applications at scale

Amazon EventBridge is a serverless event bus that makes it easier to build event-driven applications at scale using events generated from your applications, integrated Software-as-a-Service (SaaS) applications, and AWS services.

Create a new rule

Create a rule. Choose an AWS service, SaaS app or custom app as event source, define event pattern, and attach an AWS service or SaaS apps via API Destination as target(s).

Create rule

View rules

Schema registry

Find existing AWS service event schemas, generate one from an event bus, or create your own custom schemas.

How it works

▼ Getting started

Learn

Sandbox **New**

▼ Events

Event buses

Rules

Global endpoints **Now****Archives**

Replays

▼ Integration

Partner event sources

API destinations

Step 1

Define archive details

Step 2 - optional

Filter events

Amazon EventBridge > Archives

Archives

Archives are collections of events that have been published on event buses. You can create an archive here or in the [Event buses page](#). Events are continuously saved in archives, and individual events will be deleted after the retention period. An archive will persist until you manually delete it.

Archives (0/0) Loading archives



Edit

Delete

Replay

Create archive

Find archives

Any status



1



Name ▲

Status ▼

Source (Event bus) ▼

Created on ▼

Size in bytes ▼

Retention days ▼

Loading archives

Define archive details

Create an archive by defining the source and the retention period. Events are continuously saved in archives, and individual events will be deleted after the retention period. An archive will persist until you manually delete it.

Amazon EventBridge charges apply to archives. Please refer to [Amazon EventBridge Pricing](#) for details.

Archive detail

Name

Enter archive name

Maximum of 48 characters consisting of numbers, lower/upper case letters, -, _, . You can't change the name of the archive after it is created.

Step 1

Define archive details

Step 2 - optional

Filter events

Define archive details

Create an archive by defining the source and the retention period. Events are continuously saved in archives, and individual events will be deleted after the retention period. An archive will persist until you manually delete it.

Amazon EventBridge charges apply to archives. Please refer to [Amazon EventBridge Pricing](#) for details.

Archive detail

Name

archive

Maximum of 48 characters consisting of numbers, lower/upper case letters, -, _, . You can't change the name of the archive after it is created.

Archive detail

Name

archive

Maximum of 48 characters consisting of numbers, lower/upper case letters, -, _, . You can't change the name of the archive after it is created.

Description

This is a description

Maximum of 512 characters.

Source

The source, e.g. an event bus, from which the events will be archived.

default

You can't change the source of the archive after it is created.

Retention period

The number of days the archive of events will be retained. 0 is equivalent to Indefinite. The maximum is 2 billion days.

☐ Indefinite

45



days

Filter events - optional

Event pattern

Define an event pattern to filter events to be archived. If no event pattern is defined, all events from the source will be archived.

Filtering

- ☒ No event filtering
All events from the source will be archived.
- ☐ Filtering events by event pattern matching
You can define an event pattern to filter events to be archived.

Cancel

Previous

Create archive

Archives

Archives are collections of events that have been published on event buses. You can create an archive here or in the [Event buses page](#). Events are continuously saved in archives, and individual events will be deleted after the retention period. An archive will persist until you manually delete it.

Archives (1/1)



Edit

Delete

Replay

Create archive

Find archives

Any status

< 1 > ⚙

Name ▲	Status ▼	Source (Event bus) ▼	Created on ▼	Size in bytes ▼	Retention days ▼
archive	Enabled	arn:aws:events:us-east-2:082772415256:event-bus/default	Apr 29, 2022, 04:32 AM PDT	0	45 days

Rules

A rule watches for specific types of events. When a matching event occurs, the event is routed to the targets associated with the rule. A rule can be associated with one or more targets.

Select event bus

Event bus

Select or enter event bus name

default

Rules (2/2)



Edit

Delete

Enable

Create rule

Find rules

Any status

< 1 > ⚙

Name ▲	Status ▼	Type ▼	Description
Events-Archive-archive	Enabled	Managed	
admin-ec2-state-change-rule	Enabled	Standard	Run function on EC2 state change

[Amazon EventBridge](#) > [Rules](#) > [Create rule](#)

Step 1

Define rule detail

Step 2

Build event pattern

Step 3

Select target(s)

Step 4 - optional

Configure tags

Step 5

Review and create

Define rule detail [Info](#)

Rule detail

Name

rule-name

Maximum of 64 characters consisting of numbers, lower/upper case letters, -, _.

Description - optional

Enter description

Event bus [Info](#)

Select the event bus this rule applies to, either the default event bus or a custom or partner event bus.

default

☒ Enable the rule on the selected event bus

Rule type [Info](#)

☒ Rule with an event pattern

A rule that runs when an event matches the defined event pattern. EventBridge sends the event to the specified target.

☐ Schedule

A rule that runs on a schedule.

Define rule detail Info

Rule detail

Name

RunFunction

Maximum of 64 characters consisting of numbers, lower/upper case letters, `_-./`

Description - optional

Enter description

Event bus Info

Select the event bus this rule applies to, either the default event bus or a custom or partner event bus.

default

☒ Enable the rule on the selected event bus

Rule type Info



Rule with an event pattern

A rule that runs when an event matches the defined event pattern. EventBridge sends the event to the specified target.



Schedule

A rule that runs on a schedule

Amazon EventBridge > Rules > Create rule

Step 1

Define rule detail

Step 2

Define schedule

Step 3

Select target(s)

Step 4 - optional

Configure tags

Step 5

Review and create

Define schedule Info

Schedule pattern

Schedule pattern

Choose the schedule type that best meets your needs.



A fine-grained schedule that runs at a specific time, such as 8:00 a.m. PST on the first Monday of every month.



A schedule that runs at a regular rate, such as every 10 minutes.

Rate expression Info

Enter a value and the unit of time to run the schedule.

 rate ()

Value

Unit, e.g. mins, hours...

Amazon EventBridge > Rules > Create rule

Step 1

Define rule detail

Step 2

Define schedule

Step 3

Select target(s)

Step 4 - optional

Configure tags

Step 5

Review and create

Select target(s)



Permissions

Note: When using the EventBridge console, EventBridge will automatically configure the proper permissions for the selected targets. If you're using the AWS CLI, SDK, or CloudFormation, you'll need to configure the proper permissions.

Target 1

Target types

Select an EventBridge event bus, EventBridge API destination (SaaS partner), or another AWS service as a target.



EventBridge event bus



EventBridge API destination



AWS service

Select a target Info

Select target(s) to invoke when an event matches your event pattern or when schedule is triggered (limit of 5 targets per rule)

Lambda function

Function

LogEventFunction

► [Configure version/alias](#)

Function

LogEventFunction

► Configure version/alias

▼ Additional settings

Configure target input [Info](#)

You can customize the text from an event before EventBridge passes the event to the target of a rule.

Constant (JSON text)

Specify the constant in JSON

If you choose Constant (JSON text), no part of the event text is passed to the target. Instead, only the JSON text that you specify in this box is passed to the target.

```
1 {"name":"example"}
```

✓ JSON is valid

Copy

{ } Prettify JSON

Retry policy [Info](#)

A retry policy determines the maximum number of hours and number of times to retry sending an event to a target after an error occurs.

A retry policy determines the maximum number of hours and number of times to retry sending an event to a target after an error occurs.

Maximum age of event - *optional*

The maximum number of hours to keep unprocessed events for. The default value is 24 hours.

24

hour(s)

00

minute(s)

Retry attempts - *optional*

The maximum number of times to retry sending an event to a target after an error occurs. The default value is 185 times.

185

time(s)

Dead-letter queue [Info](#)

Unprocessed events can be sent to standard SQS queue

☒ None

☐ Select an Amazon SQS queue in the current AWS account to use as the dead-letter queue

☐ Specify an Amazon SQS queue in other AWS account as a dead-letter queue

Step 1
Define rule detailStep 2
Define scheduleStep 3
Select target(s)Step 4 - optional
Configure tagsStep 5
Review and createConfigure tags - *optional* [Info](#)

Tags

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

No tags associated with the resource.

[Add new tag](#)

You can add 50 more tags.

[Cancel](#)[Previous](#)[Next](#)Step 1
Define rule detailStep 2
Define scheduleStep 3
Select target(s)Step 4 - optional
Configure tagsStep 5
Review and create

Review and create

Step 1: Define rule detail

[Edit](#)

Define rule detail

Rule name
RunFunction
DescriptionStatus
✔ Enabled
Rule type
Scheduled ruleEvent bus
default


Step 2: Build schedule

[Edit](#)Event schedule [Info](#)Fixed rate of
15 minute

Step 3: Select target(s)

[Edit](#)

Targets

	Target Name	Type	Arn	Input	Role
▼	LogEventFunction 	Lambda function	 arn:aws:lambda:us-east-2:082772415256:function:LogEventFunction	Constant	-
Input to target:		Constant View			
Additional parameters:		-			
Dead-letter queue (DLQ):		-			

Step 4: Configure tag(s)

[Edit](#)

Tags

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

Key	Value

[Cancel](#)[Previous](#)[Create rule](#)

Rules

A rule watches for specific types of events. When a matching event occurs, the event is routed to the targets associated with the rule. A rule can be associated with one or more targets.

Select event bus

Event bus

Select or enter event bus name

default

Rules (3/3)

[Edit](#)[Delete](#)[Enable](#)[Create rule](#)

Any status



< 1 >



	Name	Status	Type	Description
<input type="radio"/>	Events-Archive-archive	Enabled	Managed	
<input type="radio"/>	RunFunction	Enabled	Scheduled Standard	
<input type="radio"/>	admin-ec2-state-change-rule	Enabled	Standard	Run function on EC2 state change

Step 1

Define rule detail

Step 2

Build event pattern

Step 3

Select target(s)

Step 4 - optional

Configure tags

Step 5

Review and create

Define rule detail [Info](#)

Rule detail

Name

Maximum of 64 characters consisting of numbers, lower/upper case letters, ~, -, and _.

Description - optional

Event bus [Info](#)

Select the event bus this rule applies to, either the default event bus or a custom or partner event bus.

default

☒ Enable the rule on the selected event bus

Rule type [Info](#)

☒ Rule with an event pattern

A rule that runs when an event matches the defined event pattern. EventBridge sends the event to the specified target.

☐ Schedule

A rule that runs on a schedule.

Step 1

Define rule detail

Step 2

Build event pattern

Step 3

Select target(s)

Step 4 - optional

Configure tags

Step 5

Review and create

Build event pattern [Info](#)

Event source

Event source

Select the event source from which events are sent.

☒ AWS events or EventBridge partner events

Events sent from AWS services or EventBridge partners.

☐ Other

Custom events or events sent from more than one source, e.g. events from AWS services and partners.

☐ All events

All events sent to your account.

Sample event - optional

You don't have to select or enter a sample event, but it's recommended so you can reference it when writing and testing the event pattern, or filter criteria.

You can reference the sample event when write the event pattern, or use the sample event to test if it matches the event pattern. Find a sample event, enter your own, or edit a sample event below.

Learn more about sample/test event in [Test Event Pattern](#).

☒ AWS events☐ EventBridge partner events☐ Enter my own

Event source

AWS service or EventBridge partner as source

AWS services

AWS service

The name of the AWS service as the event source

EC2

Event type

The type of events as the source of the matching pattern

EC2 Spot Instance Interruption Warning

Event pattern

Event pattern, or filter to match the events

```
1 {  
2   "source": ["aws.ec2"],  
3   "detail-type": ["EC2 Spot Instance Interruption Warning"]  
4 }
```

Copy

Test pattern

Edit pattern

Step 1

Define rule detail

Step 2

Build event pattern

Step 3

Select target(s)

Step 4 - optional

Configure tags

Step 5

Review and create

Select target(s)



Permissions

Note: When using the EventBridge console, EventBridge will automatically configure the proper permissions for the selected targets. If you're using the AWS CLI, SDK, or CloudFormation, you'll need to configure the proper permissions.

Target 1

Target types

Select an EventBridge event bus, EventBridge API destination (SaaS partner), or another AWS service as a target.

- ☐ EventBridge event bus
- ☐ EventBridge API destination
- ☒ AWS service

Select a target [Info](#)

Select target(s) to invoke when an event matches your event pattern or when schedule is triggered (limit of 5 targets per rule)

SNS topic

Topic

TopicEvents

Additional settings

Configure target input [Info](#)

You can customize the text from an event before EventBridge passes the event to the target of a rule.

Input transformer

Configure input transformer

Retry policy [Info](#)

A retry policy determines the maximum number of hours and number of times to retry sending an event to a target after an error occurs.

A retry policy determines the maximum number of hours and number of times to retry sending an event to a target after an error occurs.

Maximum age of event - optional

The maximum number of hours to keep unprocessed events for. The default value is 24 hours.

24 hour(s) 00 minute(s)

Retry attempts - optional

The maximum number of times to retry sending an event to a target after an error occurs. The default value is 185 times.

185 time(s)

Dead-letter queue [Info](#)

Unprocessed events can be sent to standard SQS queue

- ☒ None
- ☐ Select an Amazon SQS queue in the current AWS account to use as the dead-letter queue
- ☐ Specify an Amazon SQS queue in other AWS account as a dead-letter queue

Configure input transformer



 Copy

► Example input paths, Templates and Outputs

Target input transformer

You can customize the text from an event before EventBridge passes the event to the target of a rule. Using the input transformer in the console or the API, you define variables that use JSON path to reference values in the original event source. You can define up to 100 variables, assigning each a value from the input. Then you can use those variables in the Input Template as <variable-name>.

Input path

The Input Path defined as key-value pairs is used to define variables. You use JSON path to reference items in your event and store those values in variables. For instance, you could create an Input Path to reference values in the event.

```
1 {"instance": "${detail.instance-id}"}
```

Cancel

Confirm

Configure input transformer



✓ JSON is valid

Copy

Prettify JSON

Template

The Input Template is a template for the information you want to pass to your target. You can create a template that passes either a string or JSON to the target.

```
1 "The EC2 Spot Instance <instance> has been interrupted."
```

Copy

Cancel

Confirm

Amazon EventBridge > Rules > Create rule

Step 1

Define rule detail

Step 2

Build event pattern

Step 3

Select target(s)

Step 4 - optional

Configure tags

Step 5

Review and create

Configure tags - optional [Info](#)

Tags

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

No tags associated with the resource.

Add new tag

You can add 50 more tags.

Cancel

Previous

Next

Step 1
Define rule detail

Step 2
Build event pattern

Step 3
Select target(s)

Step 4 - optional
Configure tags

Step 5
Review and create

Review and create

Step 1: Define rule detail

Edit

Define rule detail

Rule name	Status	Event bus
SpotWarning	Enabled	default
Description	Rule type	
	Standard rule	

Step 2: Build event pattern

Edit

Event pattern [Info](#)

```
1 {
2   "source": ["aws-ec2"],
3   "detail-type": ["EC2 Spot Instance Interruption Warning"]
4 }
```

Copy

Step 3: Select target(s)

Edit

Targets

Target Name	Type	Arn	Input	Role
▼ TopicEvents Info	SNS topic	arn:aws:sns:us-east-2:082772415256:TopicEvents	Input transformer	-
Input to target:		Input transformer	View transformer	
Additional parameters:		-		
Dead-letter queue (DLQ):		-		

Step 4: Configure tag(s)

Edit

Tags

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value. You can use tags to search and filter your resources or track your AWS costs.

Key	Value
-----	-------

Cancel

Previous

Create rule

▼ Getting started

Learn

Sandbox [New](#)

▼ Events

Event buses

Rules

Global endpoints [New](#)

Archives

Replays

▼ Integration

Partner event sources

API destinations

▼ Schema registry

Schemas

Documentation [Info](#)

Rule SpotWarning was created successfully

Rules

A rule watches for specific types of events. When a matching event occurs, the event is routed to the targets associated with the rule. A rule can be associated with one or more targets.

Select event bus

Event bus

Select or enter event bus name

default

Rules (4/4)

Refresh

Edit

Delete

Enable

Create rule

Find rules

Any status

< 1 > Settings

	Name	Status	Type	Description
<input type="radio"/>	Events-Archive-archive	Enabled	Managed	
<input type="radio"/>	RunFunction	Enabled	Scheduled Standard	
<input type="radio"/>	SpotWarning	Enabled	Standard	
<input type="radio"/>	admin-ec2-state-change-rule	Enabled	Standard	Run function on EC2 state change

🗨️ 👤 **noahsark** Highly Voted 🏆 1 year, 1 month ago
High level notes based on the provided answers:

#6

Step 1

Go to SNS -> Topics -> Create Topic

In Type, select Standard.

In Name, input TopicEvents

Step 2

Go to EventBridge -> Rules.

In Name, input SpotWarning

In Rule Type, select Rule with an event pattern.

In Sample events, input EC2 Spot Instance Interruption Warning?

In Event Pattern, AWS Service, select EC2.

In Event Type, select EC2 Spot Instance Interruption Warning.

In Configure Target Input, select Input Transformer.

Click Configure Input Transformer.

In Sample events, select EC2 Spot Instance Interruption Warning?

In Target Input Transformer -> Input Path, enter {"instance": "\$.detail.instance-id"}

In Target Input Transformer -> Template, enter "The EC2 Spot Instance <instance> has been interrupted."

upvoted 7 times

🗨️ 👤 **smanzana** Most Recent 🔔 11 months, 1 week ago

Oh my God!! there really simulation or lab questions in this exam? We only have 140 minutes and 65 questions, which means 2 minutes per question

upvoted 1 times

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

The labs were removed from the exam.

upvoted 5 times

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

🔔 🔔 🔔 🔔.

upvoted 1 times

🗨️ 👤 **noahsark** 1 year, 1 month ago

High level notes based on the provided answers:

#4

Go to EventBridge -> Archives -> Create Archive

Specific Name not required.

Retention Period is 45 days

#5

Go to EventBridge -> Rules.

In Name, input RunFunction

In Rule Type, select Schedule.

In Occurrence, select Recurring Schedule

In Schedule type, select Rate-based schedule.

In rate, enter 15 minutes.

In Flexible time window, select Off?

In Target API, select Lambda.

In Function, look for LogEventFunction

In Payload, input {"name": "example"}

upvoted 2 times

🗨️ 👤 **helloworldabc** 1 year, 3 months ago

good~~

upvoted 1 times

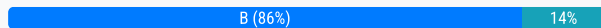
A company has a stateful, long-running workload on a single xlarge general purpose Amazon EC2 On-Demand Instance. Metrics show that the service is always using 80% of its available memory and 40% of its available CPU. A SysOps administrator must reduce the cost of the service without negatively affecting performance.

Which change in instance type will meet these requirements?

- A. Change to one large compute optimized On-Demand Instance.
- B. Change to one large memory optimized On-Demand Instance.
- C. Change to one xlarge general purpose Spot Instance.
- D. Change to two large general purpose On-Demand Instances.

Suggested Answer: B

Community vote distribution



AAAat **Highly Voted** 2 years, 3 months ago

Selected Answer: B

B because it is a stateful long running work load
upvoted 12 times

xdkonorek2 **Most Recent** 11 months, 3 weeks ago

Selected Answer: B

in general vcpu/memory(GiB) ratios for:

general purpose: 1/2

memory optimized: 1/4

compute optimized: 1/2

example:

cheapest non burstable GP: m6g.xlarge 4vpcu, 16GiB, price: 0.154

cheapest memory optimized: r6g.large 2vcpi, 16GiB, price: 0.1008

anticipated outcome: 80% cpu, 80% mem, 35% cost savings

upvoted 3 times

michaldavid 2 years ago

Selected Answer: B

BBBBBBBB

upvoted 2 times

tyfta6 2 years ago

Isn't x-Large vs large has bigger ram size?

upvoted 1 times

jipark 1 year, 4 months ago

"A memory optimized instance type would provide more memory resources."

because 80% memory already used, memory optimized is needed

upvoted 2 times

Surferbolt 2 years, 2 months ago

Selected Answer: B

B. CPU utilisation is rather low for using an xlarge instance, so a memory optimized large instance is a cheaper option that should meet the workload requirements.

upvoted 1 times

kati2k22cz 2 years, 3 months ago


Selected Answer: C

For me its C.

You can reduce teh cost using a spot instance. Here more about spot instance

<https://www.amazonaws.cn/en/ec2/spot-instances/faqs/>

upvoted 3 times

  **CiCa560** 2 years, 3 months ago

Questions states - reduce the cost of the service without negatively affecting performance.

Using a Spot Instance may affect performance, if it terminates with only a 2 min warning.

The stats state the load is Memory intensive, so B looks to be most suitable.

upvoted 7 times

A company asks a SysOps administrator to ensure that AWS CloudTrail files are not tampered with after they are created. Currently, the company uses AWS Identity and Access Management (IAM) to restrict access to specific trails. The company's security team needs the ability to trace the integrity of each file.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function when a new file is delivered. Configure the Lambda function to compute an MD5 hash check on the file and store the result in an Amazon DynamoDB table. The security team can use the values that are stored in DynamoDB to verify the integrity of the delivered files.
- B. Create an AWS Lambda function that is invoked each time a new file is delivered to the CloudTrail bucket. Configure the Lambda function to compute an MD5 hash check on the file and store the result as a tag in an Amazon S3 object. The security team can use the information in the tag to verify the integrity of the delivered files.
- C. Enable the CloudTrail file integrity feature on an Amazon S3 bucket. Create an IAM policy that grants the security team access to the file integrity logs that are stored in the S3 bucket.
- D. Enable the CloudTrail file integrity feature on the trail. The security team can use the digest file that is created by CloudTrail to verify the integrity of the delivered files.

Suggested Answer: C

Community vote distribution

D (100%)

🗳️ 👤 **Finger41** Highly Voted 3 years, 8 months ago

D - "When you enable log file integrity validation, CloudTrail creates a hash for every log file that it delivers. Every hour, CloudTrail also creates and delivers a file that references the log files for the last hour and contains a hash of each. This file is called a digest file. Validated log files are invaluable in security and forensic investigations"

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

upvoted 21 times

🗳️ 👤 **mk1523** Most Recent 6 months, 1 week ago

Selected Answer: D

DDDDDDDDDD

upvoted 1 times

🗳️ 👤 **gehadg** 8 months ago

Correct Answer: D

CloudTrail File Integrity: AWS CloudTrail offers a built-in file integrity validation feature. When this feature is enabled, CloudTrail automatically creates digest files that provide a way to verify that log files have not been tampered with. The digest files contain hashes that allow the security team to trace the integrity of each file, ensuring data has not been altered.

Operational Efficiency: This solution is the most operationally efficient, as it leverages CloudTrail's native functionality without requiring additional resources like Lambda functions or DynamoDB tables.

Other Options:

Option A and Option B: Both require custom implementations using Lambda to calculate file hashes, which increases complexity and operational overhead.

Option C: The CloudTrail file integrity feature is enabled directly on the trail itself, not just on the S3 bucket, so this option does not fully meet the requirements.

upvoted 2 times

🗳️ 👤 **2f0a02c** 1 year, 7 months ago

D is the Answer

upvoted 1 times

🗳️ 👤 **Andrew_A** 2 years ago

Selected Answer: D

AWS CloudTrail provides a feature that allows you to validate the integrity of the CloudTrail log files stored in your S3 bucket. When you enable log file integrity validation, CloudTrail creates a digest file for every log that is delivered to your account. This file contains the hash values of every log file in the delivery and can be used to confirm that the logs have not been tampered with.

upvoted 1 times

🗳️ 👤 **joesome** 2 years ago

Answer is D

upvoted 2 times

🗳️ 👤 **mavhandu** 2 years, 3 months ago

Validating CloudTrail log file integrity is important to ensure that the logs have not been tampered with and to verify their authenticity

upvoted 1 times

🗳️ 👤 **BietTuot** 2 years, 6 months ago

Selected Answer: D

D is correct answer

upvoted 1 times

🗳️ 👤 **MrMLB** 2 years, 6 months ago

Selected Answer: D

D it is

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

DDDDDDDD

upvoted 1 times

🗳️ 👤 **pravinb** 2 years, 6 months ago

only D

upvoted 1 times

🗳️ 👤 **sassy69** 2 years, 8 months ago

Selected Answer: D

I vote D as well

upvoted 1 times

🗳️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: D

D. CloudTrail log file integrity validation

upvoted 1 times

🗳️ 👤 **CHRIS12722222** 2 years, 9 months ago

integrity is not an s3 feature, it's a cloudtrail feature.

vote for D

upvoted 1 times

🗳️ 👤 **princajen** 2 years, 10 months ago

Selected Answer: D

I vote for D

upvoted 1 times

🗳️ 👤 **Mikilo** 3 years, 1 month ago

Selected Answer: D

D is the answer

upvoted 2 times

🗳️ 👤 **hammering** 3 years, 3 months ago

choose D

upvoted 1 times

When the AWS Cloud infrastructure experiences an event that may impact an organization, which AWS service can be used to see which of the organization's resources are affected?

- A. AWS Service Health Dashboard
- B. AWS Trusted Advisor
- C. AWS Personal Health Dashboard
- D. AWS Systems Manager

Suggested Answer: C

Reference:

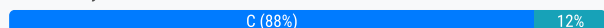
<https://docs.aws.amazon.com/health/latest/ug/getting-started-phd.html>

You can use the AWS Personal Health Dashboard to learn about AWS Health events that can affect your AWS services or account. The AWS Personal Health Dashboard presents information in two ways: a dashboard that shows recent and upcoming events organized by category, and a full event log that shows all events from the past 90 days.

To view your AWS Personal Health Dashboard

1. Sign in to the AWS Management Console and open the AWS Personal Health Dashboard at <https://phd.aws.amazon.com/phd/home>.
2. Choose **Dashboard** to view recent and upcoming events or **Event log** to view all events for the past 90 days.

Community vote distribution



umrzyj Highly Voted 3 years, 1 month ago

Selected Answer: C

"determine which of the business's resources are impacted"

PHD can provide ResourceIDs, SHD doesn't.

For every SHD there is a corresponding PHD copy but not the other way around.

upvoted 7 times

Finger41 Highly Voted 3 years, 1 month ago

Selected Answer: C

Think A and C are one in the same now - <https://aws.amazon.com/premiumsupport/technology/aws-health-dashboard/>

upvoted 5 times

gehadg Most Recent 8 months ago

Correct Answer: C

AWS Personal Health Dashboard provides a personalized view of the health of the specific AWS services and resources used by an organization. It gives real-time alerts and notifications when AWS is experiencing issues that may affect the organization's resources, allowing for targeted responses to events that might impact the infrastructure.

Other Options:

Option A (AWS Service Health Dashboard): This dashboard shows the general status of AWS services but does not provide information specific to the organization's resources.

Option B (AWS Trusted Advisor): Trusted Advisor provides recommendations on optimizing AWS resources but does not offer real-time insights into AWS infrastructure events.

Option D (AWS Systems Manager): Systems Manager helps with operational management of AWS resources but does not specifically report on AWS infrastructure events affecting resources.



upvoted 2 times

  **dante_JPMC** 1 year, 5 months ago

Selected Answer: A

The answer is A. This is talking about AWS systems being down in *GENERAL* , not your specific AWS account.

upvoted 2 times

  **jipark** 1 year, 10 months ago

Selected Answer: C

"what resources are to be affected ?" - personal Health Dashboard

upvoted 2 times

  **Mecdrex** 3 years, 2 months ago

Selected Answer: C

I vote C

upvoted 2 times

A company is using an AWS KMS customer master key (CMK) with imported key material. The company references the CMK by its alias in the Java application to encrypt data. The CMK must be rotated every 6 months.

What is the process to rotate the key?

- A. Enable automatic key rotation for the CMK, and specify a period of 6 months.
- B. Create a new CMK with new imported material, and update the key alias to point to the new CMK.
- C. Delete the current key material, and import new material into the existing CMK.
- D. Import a copy of the existing key material into a new CMK as a backup, and set the rotation schedule for 6 months.

Suggested Answer: B

Reference:

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

Community vote distribution

B (100%)

  **princajen** Highly Voted 2 years, 10 months ago

Selected Answer: B

If you choose to import keys to AWS KMS or asymmetric keys or use a custom key store, you can manually rotate them by creating a new KMS key and mapping an existing key alias from the old KMS key to the new KMS key.

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

upvoted 14 times

  **jipark** Highly Voted 1 year, 10 months ago

Selected Answer: B

why not A : Automatic key rotation is available for certain AWS-managed keys, not for imported key material.

upvoted 7 times

  **gehadg** Most Recent 8 months ago

Correct Answer: B

CMKs with Imported Key Material: For AWS KMS customer master keys (CMKs) with imported key material, automatic key rotation is not supported. To meet rotation requirements, a new CMK must be created with the updated imported key material, and the alias should be updated to point to the new CMK.

Other Options:

Option A: Automatic key rotation is not available for CMKs with imported key material.

Option C: Deleting and re-importing key material into the existing CMK does not satisfy rotation requirements, as it does not create a new CMK.

Option D: Creating a backup CMK with the same key material does not address rotation requirements and does not set up regular rotation.

upvoted 2 times

  **Christina666** 1 year, 11 months ago

Selected Answer: B

Rotation date

AWS KMS rotates key material one year (approximately 365 days) after rotation is enabled, and then every year (approximately 365 days) thereafter.

Customer managed keys

Because automatic key rotation is optional on customer managed keys and can be enabled and disabled at any time, the rotation date depends on the date that rotation was most recently enabled. That date can change many times over the life of the key.

For example, if you create a customer managed key on January 1, 2022, and enable automatic key rotation on March 15, 2022, AWS KMS rotates the key material on March 15, 2023, March 15, 2024, and every 365 days thereafter.

upvoted 1 times

  **Christina666** 1 year, 11 months ago

When you rotate KMS keys manually, you also need to update references to the KMS key ID or key ARN in your applications. Aliases, which associate a friendly name with a KMS key, can make this process easier. Use an alias to refer to a KMS key in your applications. Then, when you want to change the KMS key that the application uses, instead of editing your application code, change the target KMS key of the alias.

upvoted 1 times

  **Arnaud92** 2 years, 6 months ago

Selected Answer: B

To create new cryptographic material for your customer managed keys, you can create new KMS keys, and then change your applications or aliases to use the new KMS keys. <https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html>

upvoted 1 times

The security team is concerned because the number of AWS Identity and Access Management (IAM) policies being used in the environment is increasing. The team tasked a SysOps administrator to report on the current number of IAM policies in use and the total available IAM policies. Which AWS service should the administrator use to check how current IAM policy usage compares to current service limits?

- A. AWS Trusted Advisor
- B. Amazon Inspector
- C. AWS Config
- D. AWS Organizations


Suggested Answer: A

Reference:

<https://docs.aws.amazon.com/awssupport/latest/user/trusted-advisor-check-reference.html#iam-policies>

Community vote distribution


A (100%)

  **Christina666** Highly Voted 1 year, 11 months ago

Selected Answer: A

Key words: service limits-> Trusted Advisor

upvoted 9 times

  **Hieuly** 11 months, 1 week ago

Thanks

upvoted 1 times

  **jipark** 1 year, 10 months ago

great clue !!

upvoted 2 times

  **gehadg** Most Recent 8 months ago

Correct Answer: A

AWS Trusted Advisor provides insights into service limits and resource usage, including IAM policy limits. It offers a specific check that reports on service usage and compares it against the current service limits, allowing the administrator to see how close they are to reaching the maximum number of IAM policies.

Other Options:

Option B (Amazon Inspector): Amazon Inspector is a security assessment tool that focuses on identifying vulnerabilities in EC2 instances and applications, not on IAM policies or service limits.

Option C (AWS Config): AWS Config tracks configuration changes to AWS resources and can monitor policy compliance, but it does not provide service limit checks.

Option D (AWS Organizations): AWS Organizations is used for managing multiple AWS accounts and enabling consolidated billing, but it does not monitor service limits for IAM policies.

upvoted 3 times

  **alexiscloud** 1 year, 8 months ago

Trusted Advisor is an AWS service that provides recommendations to optimize your AWS environment. It checks your resources for:

- Security - Ensuring resources are configured securely
- Performance - Identifying underutilized resources
- Cost optimization - Ways to reduce your AWS bill
- Fault tolerance - Making your infrastructure more resilient

:A

upvoted 3 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: A

AAAAAAA

upvoted 3 times

🗉  **Surferbolt** 2 years, 8 months ago

Selected Answer: A

A. Trusted Advisor can be used to check your usage of services with service limits.

upvoted 3 times

🗉  **chanaka5** 2 years, 9 months ago

Selected Answer: A

<https://docs.aws.amazon.com/awssupport/latest/user/trusted-advisor.html>

upvoted 3 times

A SysOps administrator is trying to set up an Amazon Route 53 domain name to route traffic to a website hosted on Amazon S3. The domain name of the website is `www.example.com` and the S3 bucket name `DOC-EXAMPLE-BUCKET`. After the record set is set up in Route 53, the domain name `www.anycompany.com` does not seem to work, and the static website is not displayed in the browser.

Which of the following is a cause of this?

- A. The S3 bucket must be configured with Amazon CloudFront first.
- B. The Route 53 record set must have an IAM role that allows access to the S3 bucket.
- C. The Route 53 record set must be in the same region as the S3 bucket.
- D. The S3 bucket name must match the record set name in Route 53.

Suggested Answer: D

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/route-53-no-targets/>

Community vote distribution

D (100%)

  **princajen** Highly Voted 1 year, 10 months ago

Selected Answer: D

Check S3 website endpoint names

Make sure that the name of the resource record for your Amazon S3 website endpoint matches the name of your Amazon S3 bucket. Also, configure your bucket for website hosting.

For example, if your bucket's name is `AWSDOC-EXAMPLE-BUCKET`, the record name must also be `AWSDOC-EXAMPLE-BUCKET`.

<https://aws.amazon.com/premiumsupport/knowledge-center/route-53-no-targets/>

upvoted 16 times

  **jipark** 10 months, 3 weeks ago

tons of thanks !!

upvoted 2 times

  **michaldavid** Most Recent 1 year, 6 months ago

Selected Answer: D

DDDDDD

upvoted 2 times

  **andrerkn** 1 year, 9 months ago

Selected Answer: D

D is correct

upvoted 2 times

A SysOps administrator has used AWS CloudFormation to deploy a serverless application into a production VPC. The application consists of an AWS Lambda function, an Amazon DynamoDB table, and an Amazon API Gateway API. The SysOps administrator must delete the AWS CloudFormation stack without deleting the DynamoDB table.

Which action should the SysOps administrator take before deleting the AWS CloudFormation stack?

- A. Add a Retain deletion policy to the DynamoDB resource in the AWS CloudFormation stack.
- B. Add a Snapshot deletion policy to the DynamoDB resource in the AWS CloudFormation stack.
- C. Enable termination protection on the AWS CloudFormation stack.
- D. Update the application's IAM policy with a Deny statement for the dynamodb:DeleteTable action.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Finger41** Highly Voted 2 years, 1 month ago

Selected Answer: A

A - <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html>

upvoted 6 times

🗳️ 👤 **jipark** 10 months, 3 weeks ago

retention policy for DynamoDB

upvoted 1 times

🗳️ 👤 **Christina666** Highly Voted 11 months ago

Selected Answer: A

To keep a resource when its stack is deleted, specify Retain for that resource. You can use Retain for any resource. For example, you can retain a nested stack, Amazon S3 bucket, or EC2 instance so that you can continue to use or modify those resources after you delete their stacks.

Exception: The default policy is Snapshot for AWS::RDS::DBCluster resources and for AWS::RDS::DBInstance resources that don't specify the DBClusterIdentifier property.

upvoted 5 times

🗳️ 👤 **michaldavid** Most Recent 1 year, 6 months ago

Selected Answer: A

AAAAAA

upvoted 1 times

🗳️ 👤 **princajen** 1 year, 10 months ago

Selected Answer: A

I vote for A

upvoted 1 times

🗳️ 👤 **AWS_SJ** 2 years ago

Answer us A

upvoted 2 times

🗳️ 👤 **Mecdrex** 2 years, 2 months ago

Selected Answer: A

The answer is A

upvoted 2 times

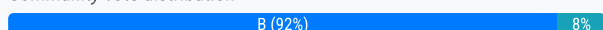
A SysOps administrator is notified that an Amazon EC2 instance has stopped responding. The AWS Management Console indicates that the system checks are failing.

What should the administrator do first to resolve this issue?

- A. Reboot the EC2 instance so it can be launched on a new host.
- B. Stop and then start the EC2 instance so that it can be launched on a new host.
- C. Terminate the EC2 instance and relaunch it.
- D. View the AWS CloudTrail log to investigate what changed on the EC2 instance.

Suggested Answer: B

Community vote distribution



🗉 👤 **lei00** Highly Voted 3 years, 9 months ago

B is correct,
system status check -- stop and restart
instance status check -- reboot

refer: https://acloud.guru/forums/aws-certified-sysops-administrator-associate/discussion/-K9I3MvysRcTvbnJHKLj/status_checks
upvoted 21 times

🗉 👤 **konieczny69** 1 year, 6 months ago

nonsense

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/monitoring-system-instance-status-check.html#system-status-checks>

system status check - start/stop or terminate

Instance status checks - reboot

upvoted 1 times

🗉 👤 **gehadg** Most Recent 8 months ago

Correct Answer: B

System Check Failures: When system checks fail on an EC2 instance, it often indicates an underlying hardware or networking issue on the physical host. Stopping and then starting the instance moves it to a different host, which usually resolves issues caused by hardware problems.

Other Options:

Option A (Reboot the EC2 instance): A reboot does not move the instance to a new host, so it will not resolve issues related to hardware failures.

Option C (Terminate and relaunch the instance): This could resolve the issue but would also result in the loss of any data stored on the instance (unless an AMI is created first). Stopping and starting is a less disruptive first step.

Option D (View the AWS CloudTrail log): CloudTrail logs API actions, but it won't provide information about hardware or networking issues affecting the instance.

upvoted 1 times

🗉 👤 **mamas_devops** 1 year, 2 months ago

B

If a system status check has failed, you can try one of the following options:

3. For an instance using an Amazon EBS-backed AMI, stop and restart the instance.

4. For an instance using an instance-store backed AMI, terminate the instance and launch a

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstances.html#InitialSteps>

upvoted 1 times

🗉 👤 **xdkonorek2** 1 year, 5 months ago

Selected Answer: B

https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/Stop_Start.html

"When you stop an instance, the instance shuts down. When you start an instance, the instance is typically migrated to a new underlying host computer"

upvoted 1 times

🗳️ 👤 **tamng** 1 year, 6 months ago

B. Stop and then start the EC2 instance so that it can be launched on a new host.

upvoted 1 times

🗳️ 👤 **konieczny69** 1 year, 6 months ago

Selected Answer: C

Correct answer is C.

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/monitoring-system-instance-status-check.html#system-status-checks>

start/stop -For instances backed by Amazon EBS, you can stop and start the instance yourself, which in most cases results in the instance being migrated to a new host. key word - may result in instance being migrated

upvoted 1 times

🗳️ 👤 **jipark** 1 year, 10 months ago

Selected Answer: B

why no A : rebooting the EC2 instance, is generally not as effective as stopping and starting it, because a reboot doesn't necessarily ensure that the instance will be launched on a new host.

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

BBBBBBB

upvoted 1 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

upvoted 1 times

🗳️ 👤 **princajen** 2 years, 10 months ago

Selected Answer: B

B is correct

upvoted 1 times

🗳️ 👤 **Finger41** 3 years, 1 month ago

Selected Answer: B

B - <https://aws.amazon.com/premiumsupport/knowledge-center/ec2-windows-system-status-check-fail/#:~:text=A%20system%20status%20check%20failure%20indicates%20a%20problem%20with%20the,from%20the%20current%20underlying%20host.>

fail/#:~:text=A%20system%20status%20check%20failure%20indicates%20a%20problem%20with%20the,from%20the%20current%20underlying%20host.

upvoted 3 times

🗳️ 👤 **Mikilo** 3 years, 1 month ago

Selected Answer: B

B is correct

upvoted 1 times

🗳️ 👤 **juraj666** 3 years, 5 months ago

Selected Answer: B

vote B - been doing this quite often lately as instances fail without prior retirement notification

upvoted 2 times

🗳️ 👤 **szl0144** 3 years, 5 months ago

vote B

upvoted 1 times

🗳️ 👤 **ngthien041292** 3 years, 7 months ago

Selected Answer: B

Vote B

upvoted 2 times

🗳️ 👤 **ahaffar** 3 years, 8 months ago

B is correct

<https://aws.amazon.com/premiumsupport/knowledge-center/ec2-windows-system-status-check-fail/>

upvoted 3 times

A software development company has multiple developers who work on the same product. Each developer must have their own development environments, and these development environments must be identical. Each development environment consists of Amazon EC2 instances and an Amazon RDS DB instance. The development environments should be created only when necessary, and they must be terminated each night to minimize costs.

What is the MOST operationally efficient solution that meets these requirements?

- A. Provide developers with access to the same AWS CloudFormation template so that they can provision their development environment when necessary. Schedule a nightly cron job on each development instance to stop all running processes to reduce CPU utilization to nearly zero.
- B. Provide developers with access to the same AWS CloudFormation template so that they can provision their development environment when necessary. Schedule a nightly Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to delete the AWS CloudFormation stacks.
- C. Provide developers with CLI commands so that they can provision their own development environment when necessary. Schedule a nightly Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to terminate all EC2 instances and the DB instance.
- D. Provide developers with CLI commands so that they can provision their own development environment when necessary. Schedule a nightly Amazon EventBridge (Amazon CloudWatch Events) rule to cause AWS CloudFormation to delete all of the development environment resources.

Suggested Answer: C

Community vote distribution

B (100%)

 **Finger41** Highly Voted 3 years, 1 month ago

Selected Answer: B

B - <https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-run-lambda-schedule.html> + <https://aws.amazon.com/blogs/mt/using-aws-lambda-to-decommission-products-provisioned-from-an-aws-service-catalog-portfolio/>

Ideally you want to delete the Stack, not just the EC2 and RDS instances. Persistent EBS volumes and snapshots may still exist that can incur cost.
upvoted 18 times

 **gehadg** Most Recent 8 months ago

Correct Answer: B


Option B provides an operationally efficient approach. By using AWS CloudFormation, each developer can easily create an identical development environment from a template. Scheduling a nightly EventBridge rule to invoke a Lambda function that deletes the CloudFormation stacks is efficient because it ensures all resources (EC2 instances, RDS instances, etc.) are automatically deleted every night, reducing costs without requiring additional manual work.

Other Options:

Option A stops processes to reduce CPU usage but does not terminate resources, meaning costs associated with instances and RDS would still accrue.

Option C and Option D require developers to manually provision environments with CLI commands, which is less efficient than CloudFormation templates. Additionally, Option C does not delete CloudFormation stacks, so it's less organized for ongoing, consistent setup and teardown.

upvoted 1 times

 **Ttomm** 1 year, 8 months ago

it should be C, keyword is "Terminate all resources", not delete stack only

upvoted 1 times

 **[Removed]** 1 year, 9 months ago

Selected Answer: B

B for sure!

upvoted 1 times

 **ronnykapo** 2 years ago

BBBBBBBBBBBB

upvoted 1 times

🗨️ 👤 **Andrew_A** 2 years ago

Selected Answer: B

This solution allows each developer to easily create an identical development environment when they need it by using the CloudFormation template. AWS CloudFormation templates are a great way to version and standardize infrastructure, and can be used to create and manage a collection of related AWS resources.

The EventBridge rule scheduled to trigger the Lambda function ensures that all resources associated with each CloudFormation stack are deleted every night, minimizing unnecessary costs.

upvoted 4 times

🗨️ 👤 **joesome** 2 years ago

B is correct

upvoted 1 times

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

Selected Answer: B

B is correct. It cannot possibly be "C" because the question stipulates the environments must be identical. Handing over access to run CLI commands will not result in identical environments, but provisioning them in CloudFormation will accomplish that.

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

BBBBBBB

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

upvoted 1 times

🗨️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: B

B is correct.

upvoted 1 times

🗨️ 👤 **princajen** 2 years, 10 months ago

Selected Answer: B

B is correct! You want to delete the stack, not just certain resources.

upvoted 1 times

🗨️ 👤 **Mikilo** 3 years, 1 month ago

Selected Answer: B

B is Correct

upvoted 1 times

🗨️ 👤 **everythingship** 3 years, 2 months ago

It's B.

upvoted 1 times

🗨️ 👤 **Mikilo** 3 years, 2 months ago

Selected Answer: B

B is Correct

upvoted 3 times

🗨️ 👤 **everythingship** 3 years, 2 months ago

Correct !

upvoted 1 times

A company is partnering with an external vendor to provide data processing services. For this integration, the vendor must host the company's data in an Amazon

S3 bucket in the vendor's AWS account. The vendor is allowing the company to provide an AWS Key Management Service (AWS KMS) key to encrypt the company's data. The vendor has provided an IAM role Amazon Resource Name (ARN) to the company for this integration. What should a SysOps administrator do to configure this integration?

- A. Create a new KMS key. Add the vendor's IAM role ARN to the KMS key policy. Provide the new KMS key ARN to the vendor.
- B. Create a new KMS key. Create a new IAM key. Add the vendor's IAM role ARN to an inline policy that is attached to the IAM user. Provide the new IAM user ARN to the vendor.
- C. Configure encryption using the KMS managed S3 key. Add the vendor's IAM role ARN to the KMS key policy. Provide the KMS managed S3 key ARN to the vendor.
- D. Configure encryption using the KMS managed S3 key. Create an S3 bucket. Add the vendor's IAM role ARN to the S3 bucket policy. Provide the S3 bucket ARN to the vendor.

Suggested Answer: D

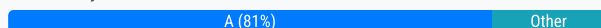
Reference:

<https://bookdown.org/bingweiliu11/aws-tutorial-book/use-case.html>

3.2 Solution

- You company's aws account (aka your personal aws account):
 - Create an admin group and an admin user
 - Create an S3 bucket with server side encryption enforced using AWS KMS service.
 - Create a KMS key to be used to encrypt files as rest
 - Ask the vendor to generate a random alpha numeric string for increased security.
 - Create an external role for the vendor
 - Attach permissions to write to the S3 bucket and use the KMS key to the external role
 - Provide the role ARN, KMS key ID and s3 bucket name to the vendor
- The vendor(cloud summit)'s AWS account:
 - Create a group for the use case
 - Identify or create users for this group
 - Attach a policy to the group to assume the role
 - User setup aws commandline tool or SDK for assume role
 - User upload files to the s3 bucket while specifying the KMS key id.

Community vote distribution



fedorian Highly Voted 2 years, 7 months ago

Selected Answer: A

The vendor is required to host the S3 bucket. It holds the company's data.

The vendor wants to use a company-provided key to encrypt the data.

So the company needs to create the new key and then provide access to that key from the IAM role which was provided by the vendor. (Answer: A)

D - Can't be D as that would mean the company is hosting the data (not the vendor). D is hosting the data at the company and providing access to the data to the vendor.

upvoted 20 times

gehadg Most Recent 8 months ago

Correct Answer: A

Option A is correct because it specifies creating a new KMS key and explicitly adding the vendor's IAM role ARN to the key policy. This approach allows the vendor to use the KMS key for encryption while ensuring access control and security through the key policy. By providing the KMS key ARN to the vendor, they can use it to encrypt the data in the S3 bucket hosted in their account.

Other Options:

Option B creates an unnecessary IAM user and an inline policy, adding complexity without directly addressing KMS encryption needs.

Option C suggests using the KMS-managed S3 key, which is controlled by AWS and does not provide the flexibility of adding external roles to the key policy.

Option D configures encryption using the KMS-managed S3 key but adds the vendor's role to the S3 bucket policy rather than the KMS key policy, which would not grant the needed access to use the key for encryption.

upvoted 1 times

🗨️ 👤 **Andrew_A** 2 years ago

Selected Answer: A

By creating a new KMS key, the SysOps administrator is ensuring that the key used to encrypt the company's data is distinct and managed separately.

The key policy is the primary resource-based policy that controls who can access and manage the key. By adding the vendor's IAM role ARN to the KMS key policy, the SysOps administrator is giving the vendor permissions to use the key, while keeping the control of the key.

By providing the ARN of the new KMS key to the vendor, the vendor will be able to use that key to encrypt the company's data stored in the S3 bucket in the vendor's account.

upvoted 2 times

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

Selected Answer: A

The provided answer links to an outside practice question - But if you go to that link ****THE QUESTION**** is different. It's as if ExamTopics has the wrong answer assigned to this question or they've pasted the wrong question into it.

As the question is written now, it's DEFINITELY ***NOT*** D. As others said - The S3 bucket needs to be in the vendor's account - So you would obviously not create one in YOUR account for this use.

upvoted 2 times

🗨️ 👤 **englishborn** 2 years, 2 months ago

Selected Answer: A

You need to create a new KMS from the question

upvoted 2 times

🗨️ 👤 **caputmundi666** 2 years, 3 months ago

Selected Answer: A

kms is in company's account. S3 is in vendor's account. Company must allow encrypt/decrypt vendor's IAM role in the KMS policy. Company should share KMS ARN of KMS.

Managed S3 KMS cannot be shared, you cannot edit its policy

upvoted 3 times

🗨️ 👤 **michele_scar** 2 years, 3 months ago

Selected Answer: A

The vendor has to host the S3, not your own company

upvoted 2 times

🗨️ 👤 **braveheart22** 2 years, 4 months ago

A is the right option from my point of view

upvoted 2 times

🗨️ 👤 **Pacoca** 2 years, 4 months ago

I agree with Fedorian

So the company needs to create the new key and then provide access to that key from the IAM role which was provided by the vendor

upvoted 1 times

🗨️ 👤 **noahsark** 2 years, 5 months ago

Selected Answer: A

<https://www.filecloud.com/supportdocs/fcdoc/latest/server/filecloud-administrator-guide/filecloud-site-setup/storage-settings/filecloud-managed-storage/s3-storage-encryption-with-aws-cross-account-kms-key>

upvoted 1 times

🗨️ 👤 **BietTuot** 2 years, 6 months ago



Selected Answer: A

I vote for A.

C. INCORRECT: You can't modify KMS managed S3 key policy.

D. INCORRECT: Because the bucket is in the vendor's account not in the company's account. Moreover, bucket policy doesn't allow Role encrypt/decrypt data. You need to use KMS Key policy.



upvoted 3 times

  **MrMLB** 2 years, 6 months ago

Selected Answer: A

A. Create a new KMS key. Add the vendor's IAM role ARN to the KMS key policy. Provide the new KMS key ARN to the vendor.


upvoted 2 times

  **tyfta6** 2 years, 6 months ago

Selected Answer: A

Vote for A

upvoted 1 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: A

Going for A

upvoted 2 times

  **Liongeek** 2 years, 7 months ago

Ans: A

upvoted 1 times

  **[Removed]** 2 years, 8 months ago

Selected Answer: A

It's A guys.

upvoted 3 times

  **zhangyu20000** 2 years, 7 months ago

question clearly ask to use KMS

upvoted 1 times

  **Surferbolt** 2 years, 8 months ago

Bucket is in vendor's account, encrypted using company's key. So the vendor will require permission to use key to access data.

upvoted 2 times

A SysOps administrator is using AWS Systems Manager Patch Manager to patch a fleet of Amazon EC2 instances. The SysOps administrator has configured a patch baseline and a maintenance window. The SysOps administrator also has used an instance tag to identify which instances to patch.

The SysOps administrator must give Systems Manager the ability to access the EC2 instances.

Which additional action must the SysOps administrator perform to meet this requirement?

- A. Add an inbound rule to the instances' security group.
- B. Attach an IAM instance profile with access to Systems Manager to the instances.
- C. Create a Systems Manager activation. Then activate the fleet of instances.
- D. Manually specify the instances to patch instead of using tag-based selection.

Suggested Answer: B

Community vote distribution

B (100%)



  **princajen** Highly Voted 10 months ago

Selected Answer: B

By default, AWS Systems Manager doesn't have permission to perform actions on your instances. Grant access by using an AWS Identity and Access Management (IAM) instance profile. An instance profile is a container that passes IAM role information to an Amazon Elastic Compute Cloud (Amazon EC2) instance at launch. You can create an instance profile for Systems Manager by attaching one or more IAM policies that define the necessary permissions to a new role or to a role you already created.

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



upvoted 14 times

  **michaldavid** Most Recent 6 months, 3 weeks ago

Selected Answer: B

bbbbbbbbb

upvoted 1 times

  **Liongeek** 7 months, 2 weeks ago

Ans: B

upvoted 1 times

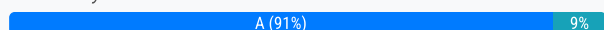
A company hosts its website on Amazon EC2 instances in the us-east-1 Region. The company is preparing to extend its website into the eu-central-1 Region, but the database must remain only in us-east-1. After deployment, the EC2 instances in eu-central-1 are unable to connect to the database in us-east-1.

What is the MOST operationally efficient solution that will resolve this connectivity issue?

- A. Create a VPC peering connection between the two Regions. Add the private IP address range of the instances to the inbound rule of the database security group.
- B. Create a VPC peering connection between the two Regions. Add the security group of the instances in eu-central-1 to the outbound rule of the database security group.
- C. Create a VPN connection between the two Regions. Add the private IP address range of the instances to the outbound rule of the database security group.
- D. Create a VPN connection between the two Regions. Add the security group of the instances in eu-central-1 to the inbound rule of the database security group.

Suggested Answer: A

Community vote distribution



princajen Highly Voted 1 year, 10 months ago

Selected Answer: A

Correct answer is A!

VPN options are out of the question.

We are left with add the IP address or a security group rule, but since you cannot create a security group rule that references a peer VPC security group, then the answer is clearly A.

upvoted 8 times

rod1234 1 year, 8 months ago

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>

upvoted 1 times

pablo23449 1 year, 8 months ago

yes, you can use SGs from peering VPNs but since it says to use in outbound the choice is A.

upvoted 4 times

caputmundi666 1 year, 3 months ago

VPC and SG are regional resources: they can't inter-operate if spread on multiple regions. So, answer is A also for this reason

upvoted 3 times

Phinx 1 year, 4 months ago

you can't peer a VPN, only VPC.

upvoted 1 times

student2020 Highly Voted 1 year, 5 months ago

A is correct.

B is wrong for 2 reasons:

a) 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.

b) it refers to outbound rule of database not the inbound rule.

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html>

upvoted 5 times

Mangesh_XI_mumbai Most Recent 6 months, 4 weeks ago

Selected Answer: A

VPC Peering and adding to inbound is key word

upvoted 2 times

Hatem08 7 months ago

Selected Answer: A

A -> initiate connection inbound


upvoted 2 times

  **callspace** 9 months ago

Selected Answer: A

This line has the answer: but the database must remain only in us-east-1. Hence us-east-1 region vpc SG need to allow the connection.

upvoted 2 times

  **jipark** 10 months, 3 weeks ago

Selected Answer: A

why not B : security groups are typically associated within the same VPC

why A : across different Regions, VPC peering is the preferred



upvoted 4 times

  **braveheart22** 1 year, 3 months ago

A is the correct answer.

B is totally wrong because adding the security group of the instances in eu-central-1 to the outbound rule of the database security group is logically adding the security group of The instances to outbound rule of database sg. Adding an INBOUND RULE to OUTBOUND RULE(outgoing traffic of the database) cannot be used to establish a VPC peering connection.

upvoted 1 times

  **MrMLB** 1 year, 6 months ago

Selected Answer: B

By creating a VPC peering connection between the two Regions and adding the security group of the instances in eu-central-1 to the outbound rule of the database security group, you can establish a direct network connection between the two VPCs and allow the instances in eu-central-1 to communicate with the database in us-east-1. This is the most operationally efficient solution because it allows for faster and more efficient communication between the two VPCs

upvoted 2 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaaaaa

upvoted 3 times

  **Liongeek** 1 year, 7 months ago

Ans: A

upvoted 1 times

A company wants to create an automated solution for all accounts managed by AWS Organizations to detect any security groups that use 0.0.0.0/0 as the source address for inbound traffic. The company also wants to automatically remediate any noncompliant security groups by restricting access to a specific CIDR block that corresponds with the company's intranet.

Which set of actions should the SysOps administrator take to create a solution?

- A. Create an AWS Config rule to detect noncompliant security groups. Set up automatic remediation to change the 0.0.0.0/0 source address to the approved CIDR block.
- B. Create an IAM policy to deny the creation of security groups that have 0.0.0.0/0 as the source address. Attach this IAM policy to every user in the company.
- C. Create an AWS Lambda function to inspect new and existing security groups. Check for a noncompliant 0.0.0.0/0 source address and change the source address to the approved CIDR block.
- D. Create a service control policy (SCP) for the organizational unit (OU) to deny the creation of security groups that have the 0.0.0.0/0 source address. Set up automatic remediation to change the 0.0.0.0/0 source address to the approved CIDR block.

Suggested Answer: A

Community vote distribution

A (100%)

  **jipark** Highly Voted 10 months, 3 weeks ago

Selected Answer: A

why A : AWS Config Rule: AWS Config rule that checks for security groups

why not D : SCPs are more suited for controlling access to AWS services and actions, not for specific security group configuration checks
upvoted 7 times

  **Mangesh_XI_mumbai** 6 months, 4 weeks ago



A is correct, related to SCPs and configuring them are not taught in SOA exam only theory is covered.

upvoted 1 times

  **Nrn143** Most Recent 1 year ago

A is the correct answer

upvoted 1 times

  **gcmrjbr** 1 year, 3 months ago

A. <https://docs.aws.amazon.com/config/latest/developerguide/vpc-sg-open-only-to-authorized-ports.html>

upvoted 2 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaa

upvoted 3 times

  **Liongeek** 1 year, 7 months ago

Ans: A



upvoted 2 times

  **Surferbolt** 1 year, 8 months ago

Selected Answer: A

A. It's a job for Config.

upvoted 3 times

  **Rick365** 1 year, 10 months ago

Selected Answer: A

A. Create an AWS Config rule to detect noncompliant security groups. Set up automatic remediation to change the 0.0.0.0/0 source address to the approved CIDR block?

upvoted 2 times

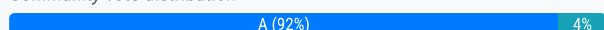
A company requires that all activity in its AWS account be logged using AWS CloudTrail. Additionally, a SysOps administrator must know when CloudTrail log files are modified or deleted.

How should the SysOps administrator meet these requirements?

- A. Enable log file integrity validation. Use the AWS CLI to validate the log files.
- B. Enable log file integrity validation. Use the AWS CloudTrail Processing Library to validate the log files.
- C. Use CloudTrail Insights to monitor the log files for modifications.
- D. Use Amazon CloudWatch Logs to monitor the log files for modifications.

Suggested Answer: B

Community vote distribution



Andrew_A Highly Voted 2 years ago

Selected Answer: A

Option B is incorrect because AWS CloudTrail Processing Library helps developers to read, process, and analyze AWS CloudTrail data but doesn't provide the functionality to validate the integrity of CloudTrail log files.

upvoted 6 times

XXXXXINN Most Recent 7 months, 3 weeks ago

Why no one select D?

upvoted 1 times

pekalyok 1 year, 2 months ago

Selected Answer: D

While the other options have their uses, they don't directly meet the requirement as effectively as option D:

A and B (Log File Integrity Validation): Enabling log file integrity validation is important for ensuring that the logs have not been tampered with. However, this feature is more about post-event validation rather than real-time monitoring or alerting. It requires manual initiation (using the AWS CLI or CloudTrail Processing Library) to validate the integrity of log files, which does not provide immediate notifications of modifications or deletions.

C (CloudTrail Insights): CloudTrail Insights is designed to identify unusual operational activity within your AWS account, not specifically to monitor log file integrity or alert on log file modifications or deletions. It is more focused on detecting anomalous API activity rather than changes to the log files themselves.

upvoted 1 times

McEgowan2023 1 year, 7 months ago

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them.

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

upvoted 2 times

Christina666 1 year, 11 months ago

Selected Answer: A

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-cli.html>

To validate logs with the AWS Command Line Interface, use the CloudTrail validate-logs command. The command uses the digest files delivered to your Amazon S3 bucket to perform the validation. For information about digest files, see CloudTrail digest file structure.

The AWS CLI allows you to detect the following types of changes:

Modification or deletion of CloudTrail log files

Modification or deletion of CloudTrail digest files

Modification or deletion of both of the above

upvoted 2 times

🗨️ **braveheart22** 2 years, 3 months ago

AAAA is the correct answer.

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them.

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

upvoted 2 times

🗨️ **braveheart22** 2 years, 4 months ago

I agree with foreverlearner, the correct answer is AAAAA.

upvoted 1 times

🗨️ **zolthar_z** 2 years, 6 months ago

Selected Answer: A

The answer is A, the cloud trail processing library is only to process logs, not check integrity

upvoted 3 times

🗨️ **MrMLB** 2 years, 6 months ago

Selected Answer: B

B

Option A is incorrect because it does not specify how to validate the log files. Option C is incorrect because CloudTrail Insights is a feature that allows you to analyze CloudTrail log data, but it does not provide a way to validate log file integrity. Option D is incorrect because Amazon CloudWatch Logs is a service that allows you to monitor, store, and access your log data, but it does not provide a way to validate log file integrity.

upvoted 1 times

🗨️ **foreverlearner** 2 years, 6 months ago

Another wrong ChatGPT answer.. "To validate the integrity of CloudTrail log files, you can use the AWS CLI or create your own solution" (<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>)

upvoted 6 times

🗨️ **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaa

upvoted 2 times

🗨️ **Surferbolt** 2 years, 8 months ago

A is the answer.

upvoted 2 times

🗨️ **AAAaat** 2 years, 9 months ago

Selected Answer: A

Yes it is A

upvoted 3 times

🗨️ **haxaffee** 2 years, 9 months ago

Selected Answer: A

Answer can only be A. How to use CLI -> <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-cli.html>

upvoted 3 times

🗨️ **princajen** 2 years, 10 months ago

Selected Answer: A

The answer is A!

The CloudTrail Processing Library is a Java library that provides an easy way to process AWS CloudTrail logs. You provide configuration details about your CloudTrail SQS queue and write code to process events. The CloudTrail Processing Library does the rest. It polls your Amazon SQS queue, reads and parses queue messages, downloads CloudTrail log files, parses events in the log files, and passes the events to your code as Java objects.

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/use-the-cloudtrail-processing-library.html>

upvoted 3 times

  **[Removed]** 2 years, 6 months ago

princajen, your response describes answer B.

upvoted 1 times

  **Flosuccess** 2 years, 10 months ago

Looks like the answer is A

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

upvoted 2 times

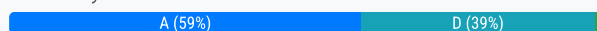
A company is planning to host its stateful web-based applications on AWS. A SysOps administrator is using an Auto Scaling group of Amazon EC2 instances. The web applications will run 24 hours a day, 7 days a week throughout the year. The company must be able to change the instance type within the same instance family later in the year based on the traffic and usage patterns.

Which EC2 instance purchasing option will meet these requirements MOST cost-effectively?

- A. Convertible Reserved Instances
- B. On-Demand Instances
- C. Spot Instances
- D. Standard Reserved Instances

Suggested Answer: A

Community vote distribution



foreverlearner Highly Voted 2 years, 6 months ago

Selected Answer: A

The question says: "must be able to change the instance TYPE within the SAME instance FAMILY". Note that the question says TYPE, and not SIZE. According to <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ri-modifying.html>, You can "Change the instance SIZE within the same instance family and generation" and

"For example, you can modify a Reserved Instance from t2.small to t2.large because they're both in the same T2 family and generation. But you can't modify a Reserved Instance from T2 to M2 or from T2 to T3, because in both these examples, the target instance family and generation are not the same as the original Reserved Instance."

In the above example: General Purpose Family, TYPE T, Generation 2, SIZE small (T2.large)

(<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html#AvailableInstanceTypes>)

Convert it to T2.large. This is same TYPE (T2), different SIZE (large). Hence, Standard would cover it.

However, if I want to change T3.small, this would be a different TYPE (T3) within the same family (General Purpose).

Hope this example helps clarifying.

upvoted 20 times

kraytom 1 year, 11 months ago

This is wrong. Visit the linked userguide on instance types and see for yourself. The T is the family. Instance type refers to the entire t2.large. large is the size. the other comment correcting this is also wrong.

upvoted 3 times

vivanchyk 1 year, 3 months ago

let's define what instance family, type (and size) mean, shall we

e.g. "t2.small" instance

- instance family "t2"
- instance size ".small"
- instance type "t2.small"

so if the question requires the ability to change "the instance type within the same instance family", it means persistence of "t2" part and changing anything that goes after.

upvoted 3 times

henryford 2 years, 1 month ago

This is wrong. TYPE and SIZE are the same thing in this context. What you're saying to be the "TYPE" is actually the FAMILY. t2.large = Family T, Generation 2, Type Large. General Purpose is not the family.

upvoted 6 times

AminTriton Most Recent 3 months, 2 weeks ago

Selected Answer: A

To those who answered D: Standard Reserved Instance has a fixed Family but is flexible in Size. The question asks for changing type (which means family as well) therefore D cannot be the answer. Check the definition of Family, Size, and Type in this link <https://docs.aws.amazon.com/ec2/latest/instancetypes/ec2-types.pdf> on page 4

upvoted 1 times

🗨️ **ede83d8** 4 months, 3 weeks ago

Selected Answer: D

Enables you to modify Availability Zone, scope, networking type, and instance size (within the same instance family) of your Reserved Instance. For more information, see Modifying Reserved Instances.

upvoted 1 times

🗨️ **OlehKom** 6 months, 2 weeks ago

Selected Answer: A

"..must be able to change the instance type within the same instance family later in the year "

upvoted 1 times

🗨️ **numark** 7 months, 1 week ago

Answer ia A and ChatGPT agrees

Solution: Convertible Reserved Instances

Cost-Effectiveness: Reserved Instances (RIs) offer significant discounts (up to 72%) compared to On-Demand pricing when the instances run continuously.

Flexibility: Convertible RIs allow you to change the instance type, size, or even the operating system within the same family, which matches the company's requirement to adjust instance types later based on traffic. Commitment Period: Reserved Instances require a 1-year or 3-year commitment, suitable for 24/7 workloads.

upvoted 1 times

🗨️ **numark** 6 months, 3 weeks ago

Standard RIs do not offer the same level of flexibility to change instance types once purchased. If the company is confident that there will be no need to change instance types or attributes, standard RIs could potentially be more cost-effective. But, since the company anticipates needing to adjust instance types based on traffic patterns, this option is less suitable.

upvoted 1 times

🗨️ **Aamee** 9 months, 3 weeks ago

BC pooree awaam sirf A aur D hee main larr rahee hai yahan and still there's no final consensus on it! :(

upvoted 1 times

🗨️ **acnaz** 11 months, 1 week ago

By GTP4o

Given the company's need to run applications continuously throughout the year with the ability to adapt instance types based on traffic, Convertible Reserved Instances provide the best balance of cost savings and flexibility. This purchasing option ensures the company can optimize costs while retaining the flexibility to adjust their infrastructure as usage patterns evolve.

upvoted 1 times

🗨️ **NSA_Poker** 11 months ago

I tell my buddy this all the time, "ChatGPT will NOT teach you AWS!"

upvoted 1 times

🗨️ **auxwww** 1 year ago

Selected Answer: D

Change Type within the Instance's Family - Standard RIs allow this and they are cheaper. You don't require convertible RIs! period.

upvoted 1 times

🗨️ **64e0ca8** 1 year ago

Convertible Reserved Instances allow for updating the instance type

upvoted 1 times

🗨️ **mamas_devops** 1 year, 2 months ago

Depending on context AWS uses both terms TYPE and SIZE to point on instance NAME (eg t3.medium). That's why many people got confused.

Here it is SIZE: <https://aws.amazon.com/ec2/instance-types/>

Here it is TYPE: <https://docs.aws.amazon.com/ec2/latest/instancetypes/gp.html>

In context of RI it states the following:

Instance type: For example, m4.large. This is composed of the instance family (for example, m4) and the instance size (for example, large).

In question the keys are 'within the same instance family' and 'most cost-effective'

Standard RI allows to change instance type within the same family (m4.large -> m4.xlarge) and it is cheaper then Convertible RI

upvoted 1 times

🗨️ **vivanchyk** 1 year, 3 months ago

Selected Answer: D

ok, the final comment to solve the argument in here

as mentioned previously, the main clause in the question is "The company must be able to change __the instance type within the same instance family later__"

so let's define what instance family, type (and size) mean, shall we

e.g. "t2.small" instance

- instance family "t2"

- instance size ".small"

- instance type "t2.small"

so if the question requires the ability to change "the instance type within the same instance family", it means persistence of "t2" part and changing anything that goes after.

upvoted 4 times

🗨️ **stoy123** 1 year, 4 months ago

Selected Answer: A

A!!!

C is not correct because:

Standard Reserved Instance enables you to modify ... instance size within the same instance type

upvoted 2 times

🗨️ **nosense** 1 year, 4 months ago

Selected Answer: D

d

Convertible Reserved Instances (up to 55%) Standard Reserved Instances (up to 75%)

upvoted 1 times

🗨️ **Suksay** 1 year, 4 months ago

Selected Answer: D

D because the text say "able to change the instance type within the same instance family" (type mean size here) and even if it's possible to change instance size and type with Convertible Reserved Instances (up to 55%) Standard Reserved Instances (up to 75%) is more cheaper than Convertible Reserved Instances

upvoted 2 times

🗨️ **xdkonorek2** 1 year, 5 months ago

Selected Answer: A

I will say A because there is a need to change instance type within instance family. With D you wouldn't be able to change instance generation and it's part of instance type.

upvoted 2 times

🗨️ **tamng** 1 year, 6 months ago

A. Convertible Reserved Instances

upvoted 1 times

🗨️ **tamng** 1 year, 6 months ago

"The company must be able to change the instance type within the same instance family later in the year based on the traffic and usage patterns."

=> A

upvoted 1 times

🗨️ **konieczny69** 1 year, 6 months ago

Selected Answer: D

<https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/standard-vs.-convertible-offering-classes.html>

instance type contains family

reserved are cheaper

Convertible Reserved Instances are useful when:

Purchasing Reserved Instances in the payer account instead of a subaccount. You can more easily modify Convertible Reserved Instances to meet changing needs across your organization.

Workloads are likely to change. In this case, a Convertible Reserved Instance enables you to adapt as needs evolve while still obtaining discounts and capacity reservations.

You want to hedge against possible future price drops.

You can't or don't want to ask teams to do capacity planning or forecasting.

You expect compute usage to remain at the committed amount over the commitment period.

upvoted 1 times

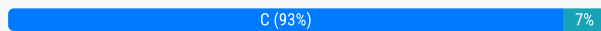
An application runs on Amazon EC2 instances in an Auto Scaling group. Following the deployment of a new feature on the EC2 instances, some instances were marked as unhealthy and then replaced by the Auto Scaling group. The EC2 instances terminated before a SysOps administrator could determine the cause of the health status changes. To troubleshoot this issue, the SysOps administrator wants to ensure that an AWS Lambda function is invoked in this situation.

How should the SysOps administrator meet these requirements?

- A. Activate the instance scale-in protection setting for the Auto Scaling group. Invoke the Lambda function through Amazon EventBridge (Amazon CloudWatch Events).
- B. Activate the instance scale-in protection setting for the Auto Scaling group. Invoke the Lambda function through Amazon Route 53.
- C. Add a lifecycle hook to the Auto Scaling group to invoke the Lambda function through Amazon EventBridge (Amazon CloudWatch Events).
- D. Add a lifecycle hook to the Auto Scaling group to invoke the Lambda function through Amazon Route 53.

Suggested Answer: C

Community vote distribution



Atown Highly Voted 2 years, 1 month ago

Selected Answer: C

A is Wrong because Instance scale-in protection does not protect Auto Scaling instances from the following: Health check replacement if the instance fails health checks

C is the Correct Answers -- When a scale-in event occurs, a lifecycle hook pauses the instance before it is terminated and sends you a notification using Amazon EventBridge. While the instance is in the wait state, you can invoke an AWS Lambda function or connect to the instance to download logs or other data before the instance is fully terminated.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

upvoted 17 times

jipark 1 year, 4 months ago

I choose C

upvoted 1 times

kati2k22cz Highly Voted 2 years, 3 months ago

Selected Answer: C

OK its C.

check this doc <https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

upvoted 6 times

braveheart22 Most Recent 1 year, 10 months ago

Option C is the way

upvoted 3 times

ahalamri 1 year, 9 months ago

This is the way.

upvoted 1 times

michaldavid 2 years ago

Selected Answer: C

ccccccc

upvoted 2 times

Liongeek 2 years, 1 month ago

Ans: C

upvoted 2 times

bakjeeone 2 years, 2 months ago

Selected Answer: C

Answer is C

upvoted 4 times

  **kati2k22cz** 2 years, 3 months ago

Selected Answer: A

think better. Its A

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

upvoted 2 times

  **tamng** 1 year ago

C not A

upvoted 3 times

  **pablo23449** 2 years, 2 months ago

Instance scale-in protection does not protect Auto Scaling instances from the following: Health check replacement if the instance fails health checks

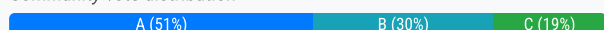
upvoted 2 times

A company runs an application that hosts critical data for several clients. The company uses AWS CloudTrail to track user activities on various AWS resources. To meet new security requirements, the company needs to protect the CloudTrail log files from being modified, deleted, or forged. Which solution will meet these requirements?

- A. Enable CloudTrail log file integrity validation.
- B. Use Amazon S3 MFA Delete on the S3 bucket where the CloudTrail log files are stored.
- C. Use Amazon S3 Versioning to keep all versions of the CloudTrail log files.
- D. Use AWS Key Management Service (AWS KMS) security keys to secure the CloudTrail log files.

Suggested Answer: C

Community vote distribution



🗳️ **AminTriton** 3 months, 2 weeks ago

Selected Answer: A

If you have answered anything but A, remember that CloudTrail log file integrity validation is EXACTLY made for this purpose and nothing else.
upvoted 1 times

🗳️ **wakburn** 6 months, 3 weeks ago

Selected Answer: A

A. Enable CloudTrail log file integrity validation.

Enabling CloudTrail log file integrity validation ensures that the log files are protected using industry-standard algorithms (SHA-256 for hashing and SHA-256 with RSA for digital signing). This makes it computationally infeasible to modify, delete, or forge the log files without detection[1] (<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>)[2] (<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/best-practices-security.html>).
upvoted 3 times

🗳️ **jagdishmav** 7 months ago

Selected Answer: A

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them.

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html#cloudtrail-log-file-validation-intro-use-cases>
upvoted 1 times

🗳️ **numark** 7 months, 1 week ago

Answer is A and ChatGPT agrees:

A. Enable CloudTrail log file integrity validation (Correct)

CloudTrail log file integrity validation uses SHA-256 hashing and digital signatures to ensure that log files have not been tampered with or altered. This provides cryptographic proof of the integrity of the log files. It directly addresses the requirement to ensure that logs are not forged or modified. This is the most relevant feature provided by AWS CloudTrail for log protection.

B. Use Amazon S3 MFA Delete on the S3 bucket where the CloudTrail log files are stored (Partially Correct) MFA Delete prevents accidental or unauthorized deletions of S3 objects, but it does not protect against forgery or tampering of log contents. While useful, this alone does not meet the full requirement of securing logs against modification or forgery.

upvoted 1 times

🗳️ **MintTeaClarity** 7 months, 3 weeks ago

Selected Answer: B

I think it's B.

upvoted 1 times

🗳️ **Aamee** 9 months, 3 weeks ago

Selected Answer: B

I'd go with B here as it specifically asks to protect the CT files from being modified, deleted or forged.

upvoted 2 times

🗳️ 👤 **rcptryk** 1 year ago

Selected Answer: A

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html#:~:text=To%20determine%20whether,CloudTrail%20delivered%20them>.

upvoted 1 times

🗳️ 👤 **stoy123** 1 year, 4 months ago

no single option among the provided choices completely fulfills all three requirements

upvoted 2 times

🗳️ 👤 **Invr11** 1 year, 4 months ago

Selected Answer: B

"Enable MFA delete on the Amazon S3 bucket where you store log files

When you configure multi-factor authentication (MFA), attempts to change the versioning state of bucket, or delete an object version in a bucket, require additional authentication. This way, even if a user acquires the password of an IAM user with permissions to permanently delete Amazon S3 objects, you can still prevent operations that could compromise your log files."

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/best-practices-security.html>

upvoted 4 times

🗳️ 👤 **dante_JPMC** 1 year, 5 months ago

Selected Answer: A

There's no question at all it's A.

upvoted 4 times

🗳️ 👤 **kret** 1 year, 3 months ago

There's not question A would not PROTECT from those action. It will give an option to detect those actions.

upvoted 1 times

🗳️ 👤 **Rabbit117** 1 year, 5 months ago

Selected Answer: B

I think B is correct. CloudTrail Log integrity does not prevent the file from being deleted. It will only detect if the file was deleted, to prevent deletion you should use MFA delete, which requires versioning to be enabled.

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

upvoted 4 times

🗳️ 👤 **ogogundare** 1 year, 5 months ago

Selected Answer: B

B is the correct answer. The key word here is to prevent the logs from being deleted.

Answer would be incorrect as it does not mention about tracking who deleted the file.

upvoted 3 times

🗳️ 👤 **tamng** 1 year, 6 months ago

A is correct

upvoted 2 times

🗳️ 👤 **Saibal9** 1 year, 6 months ago

Answer b is correct. We not only don't want the files to be tampered with, no one ought to be able to delete it either.

So, the only option that looks correct is option b.

upvoted 1 times

🗳️ 👤 **tamng** 1 year, 6 months ago

you wrong, A is correct not B

upvoted 1 times

🗳️ 👤 **konieczny69** 1 year, 6 months ago

Selected Answer: B

real answer is WORM, aka object locks: <https://aws.amazon.com/blogs/storage/protecting-data-with-amazon-s3-object-lock/> - but its not in the answers so we need to pick from what we have

A is incorrect - <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html> - integrity feature doe not

prevent anyone from deleting the file

D is nonsense

so its either B or C

since MFA delete inherently requires versioning, I'd go for B

upvoted 3 times

🗨️ 👤 **konieczny69** 1 year, 6 months ago

ChatGTP suggested A in the first place. I replied with a piece of documentation about log file validation it switched answer to B

upvoted 3 times

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

A is the right answer, how C, there is no mentioning of s3, it is mentioned "on various resource". i oppose C

upvoted 1 times

🗨️ 👤 **Hatem08** 1 year, 7 months ago

Selected Answer: C

Def. C is the correct as versioning on s3 buckets keeps the objects safe from deletion, while A won't prevent the deletion action itself it will just show it...

upvoted 1 times

A global company operates out of five AWS Regions. A SysOps administrator wants to identify all the company's tagged and untagged Amazon EC2 instances.

The company requires the output to display the instance ID and tags.



What is the MOST operationally efficient way for the SysOps administrator to meet these requirements?

- A. Create a tag-based resource group in AWS Resource Groups.
- B. Use AWS Trusted Advisor. Export the EC2 On-Demand Instances check results from Trusted Advisor.
- C. Use Cost Explorer. Choose a service type of EC2-Instances, and group by Resource.
- D. Use Tag Editor in AWS Resource Groups. Select all Regions, and choose a resource type of AWS::EC2::Instance.

Suggested Answer: D

Community vote distribution

D (100%)



  **princajen**  10 months ago

Selected Answer: D

With Tag Editor, you build a query to find resources in one or more AWS Regions that are available for tagging. You can choose up to 20 individual resource types, or build a query on All resource types. Your query can include resources that already have tags, or resources that have no tags.

<https://docs.aws.amazon.com/ARG/latest/userguide/find-resources-to-tag.html>



upvoted 15 times

  **michaldavid**  6 months, 3 weeks ago

Selected Answer: D

ddddddd

upvoted 4 times

  **Liongeek** 7 months, 2 weeks ago

You can't create a tag-based group if the instances may not be tagged --

Ans: D

upvoted 3 times

A company needs to upload gigabytes of files every day. The company need to achieve higher throughput and upload speeds to Amazon S3. Which action should a SysOps administrator take to meet this requirement?

- A. Create an Amazon CloudFront distribution with the GET HTTP method allowed and the S3 bucket as an origin.
- B. Create an Amazon ElastiCache cluster and enable caching for the S3 bucket.
- C. Set up AWS Global Accelerator and configure it with the S3 bucket.
- D. Enable S3 Transfer Acceleration and use the acceleration endpoint when uploading files.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **jipark** Highly Voted 1 year, 10 months ago

Selected Answer: D

why not C :

C. AWS Global Accelerator is used to improve the availability and performance of applications for globally distributed users, but it doesn't directly impact upload speeds to S3.

upvoted 11 times

🗳️ **princajen** Highly Voted 2 years, 10 months ago

Selected Answer: D

I vote for D!

Enable Amazon S3 Transfer Acceleration

Amazon S3 Transfer Acceleration can provide fast and secure transfers over long distances between your client and Amazon S3. Transfer Acceleration uses Amazon CloudFront's globally distributed edge locations.

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/>

upvoted 6 times

🗳️ **jagdishmav** Most Recent 7 months ago

Selected Answer: D

Activate Amazon S3 Transfer Acceleration

Amazon S3 Transfer Acceleration provides fast and secure transfers over long distances between your client and Amazon S3. Transfer Acceleration uses the globally distributed edge locations of Amazon CloudFront.

Transfer Acceleration incurs additional charges, so be sure that you review pricing. To determine if Transfer Acceleration improves the transfer speeds for your use case, review the Amazon S3 Transfer Acceleration Speed Comparison tool.

upvoted 1 times

🗳️ **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddd

upvoted 1 times

🗳️ **Liongeek** 2 years, 7 months ago

Ans: D

upvoted 1 times

🗳️ **Atown** 2 years, 7 months ago

Selected Answer: D

D is the correct answer

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/>

Enable Amazon S3 Transfer Acceleration

Amazon S3 Transfer Acceleration can provide fast and secure transfers over long distances between your client and Amazon S3. Transfer

Acceleration uses Amazon CloudFront's globally distributed edge locations.

Transfer Acceleration incurs additional charges, so be sure to review pricing. To determine if Transfer Acceleration might improve the transfer speeds for your use case, review the Amazon S3 Transfer Acceleration Speed Comparison tool.

Note: Transfer Acceleration doesn't support cross-Region copies using CopyObject.

upvoted 2 times

  **Surferbolt** 2 years, 8 months ago

Selected Answer: D

D is the answer.

upvoted 1 times

A SysOps administrator maintains the security and compliance of a company's AWS account. To ensure the company's Amazon EC2 instances are following company policy, a SysOps administrator wants to terminate any EC2 instance that do not contain a department tag. Noncompliant resources must be terminated in near-real time.

Which solution will meet these requirements?

- A. Create an AWS Config rule with the required-tags managed rule to identify noncompliant resources. Configure automatic remediation to run the AWS- TerminateEC2Instance automation document to terminate noncompliant resources.
- B. Create a new Amazon EventBridge (Amazon CloudWatch Events) rule to monitor when new EC2 instances are created. Send the event to a Simple Notification Service (Amazon SNS) topic for automatic remediation.
- C. Ensure all users who can create EC2 instances also have the permissions to use the ec2:CreateTags and ec2:DescribeTags actions. Change the instance's shutdown behavior to terminate.
- D. Ensure AWS Systems Manager Compliance is configured to manage the EC2 instances. Call the AWS-StopEC2Instances automation document to stop noncompliant resources.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **kati2k22cz** Highly Voted 1 year, 9 months ago

Selected Answer: A

vote for A.

https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_use-managed-rules.html

upvoted 6 times

🗳️ **jjpark** Most Recent 10 months, 3 weeks ago

Selected Answer: A

config rule can protect many behavior before done.

upvoted 4 times

🗳️ **DZRomero** 11 months, 1 week ago

Selected Answer: A

A.

D IS NOT AN OPTION

upvoted 2 times

🗳️ **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaaa

upvoted 4 times

🗳️ **Liongeek** 1 year, 7 months ago

Ans: A

upvoted 2 times

🗳️ **Surferbolt** 1 year, 8 months ago

Selected Answer: A

A. Config.

upvoted 4 times

A company uploaded its website files to an Amazon S3 bucket that has S3 Versioning enabled. The company uses an Amazon CloudFront distribution with the S3 bucket as the origin. The company recently modified the files, but the object names remained the same. Users report that old content is still appearing on the website.

How should a SysOps administrator remediate this issue?

- A. Create a CloudFront invalidation, and add the path of the updated files.
- B. Create a CloudFront signed URL to update each object immediately.
- C. Configure an S3 origin access identity (OAI) to display only the updated files to users.
- D. Disable S3 Versioning on the S3 bucket so that the updated files can replace the old files.

Suggested Answer: A

Community vote distribution

A (100%)

  **princajen** Highly Voted 10 months ago

Selected Answer: A

Answer is A!

.Short description

By default, CloudFront caches a response from Amazon S3 for 24 hours (Default TTL of 86,400 seconds). If your request lands at an edge location that served the Amazon S3 response within 24 hours, then CloudFront uses the cached response. This happens even if you updated the content in Amazon S3.



Use one of the following ways to push the updated Amazon S3 content from CloudFront:

Invalidate the Amazon S3 objects.

Use object versioning

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-serving-outdated-content-s3/>



upvoted 13 times

  **michaldavid** Most Recent 6 months, 3 weeks ago

Selected Answer: A

aaaaaa

upvoted 2 times

  **Liongeek** 7 months, 2 weeks ago

Ans: A

upvoted 2 times

  **Surferbolt** 8 months, 2 weeks ago

Selected Answer: A

A is the answer.

upvoted 2 times

A company has two VPC networks named VPC A and VPC B. The VPC A CIDR block is 10.0.0.0/16 and the VPC B CIDR block is 172.31.0.0/16. The company wants to establish a VPC peering connection named pcx-12345 between both VPCs.


Which rules should appear in the route table of VPC A after configuration? (Choose two.)

- A. Destination: 10.0.0.0/16, Target: Local
- B. Destination: 172.31.0.0/16, Target: Local
- C. Destination: 10.0.0.0/16, Target: pcx-12345
- D. Destination: 172.31.0.0/16, Target: pcx-12345
- E. Destination: 10.0.0.0/16, Target: 172.31.0.0/16

Suggested Answer: AD

Community vote distribution

AD (100%)

 **kati2k22cz** Highly Voted 2 years, 9 months ago


Selected Answer: AD

Yes, A and D.

lookthe table on this doc

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-routing.html>

upvoted 10 times

 **be9z** Most Recent 12 months ago

Local Route:

Destination: VPC A CIDR (10.0.0.0/16)

Target: Local (implicitly points to the VPC itself)

Explanation: This route ensures that traffic within VPC A (between instances with IP addresses in the 10.0.0.0/16 range) stays local and doesn't need to go through the peering connection.

Route to Peer VPC B:

Destination: VPC B CIDR (172.31.0.0/16)

Target: VPC peering connection ID (pcx-12345)

Explanation: This route directs traffic destined for VPC B (IP addresses in the 172.31.0.0/16 range) to the VPC peering connection, allowing communication between instances in VPC A and VPC B.

upvoted 3 times


 **piavik** 2 years, 2 months ago

Selected Answer: AD

While AD is correct, the question itself is not quite clear.

Option A - routes already are there, they not "will appear" after configuring VPC peering

upvoted 2 times

 **braveheart22** 2 years, 3 months ago

A and D are the right answers

upvoted 2 times

 **michaldavid** 2 years, 6 months ago

Selected Answer: AD

A and D

upvoted 2 times

 **Liongeek** 2 years, 7 months ago

Ans: D

upvoted 1 times

 **Atown** 2 years, 7 months ago

Selected Answer: AD

I agree with A and D

upvoted 2 times

A company analyzes sales data for its customers. Customers upload files to one of the company's Amazon S3 buckets, and a message is posted to an Amazon Simple Queue Service (Amazon SQS) queue that contains the object Amazon Resource Name (ARN). An application that runs on an Amazon EC2 instance polls the queue and processes the messages. The processing time depends on the size of the file. Customers are reporting delays in the processing of their files. A SysOps administrator decides to configure Amazon EC2 Auto Scaling as the first step. The

SysOps administrator creates an Amazon Machine Image (AMI) that is based on the existing EC2 instance. The SysOps administrator also creates a launch template that references the AMI.

How should the SysOps administrator configure the Auto Scaling policy to improve the response time?

- A. Add several different instance sizes in the launch template. Create an Auto Scaling policy based on the ApproximateNumberOfMessagesVisible metric to select the size of the instance based on the number of messages in the queue.
- B. Create an Auto Scaling policy based on the ApproximateNumberOfMessagesDelayed metric to scale the number of instances based on the number of messages in the queue that have been delayed.
- C. Create a custom metric based on the ASGAverageCPUUtilization metric and the GroupPendingInstances metric from the Auto Scaling group. Modify the application to calculate the metric and post the metric to Amazon CloudWatch once each minute. Create an Auto Scaling policy based on this metric to scale the number of instances.
- D. Create a custom metric based on the ApproximateNumberOfMessagesVisible metric and the number of instances in the InService state in the Auto Scaling group. Modify the application to calculate the metric and post the metric to Amazon CloudWatch once each minute. Create an Auto Scaling policy based on this metric to scale the number of instances.

Suggested Answer: B

Community vote distribution

D (81%)

B (19%)

  **numark** 7 months, 1 week ago

It is D: ApproximateNumberOfMessagesVisible Metric:

This metric indicates the number of messages waiting to be processed in the Amazon SQS queue. Number of InService Instances: The scaling decision should consider the number of currently available EC2 instances in the Auto Scaling group to handle the load efficiently.

Custom Metric: By creating a custom metric that divides the ApproximateNumberOfMessagesVisible by the number of running (InService) instances, you can calculate the per-instance workload. Posting the Metric to CloudWatch: Once the application posts this custom metric to Amazon CloudWatch, you can create an Auto Scaling policy that uses it to scale the number of instances dynamically.

B: ApproximateNumberOfMessagesDelayed refers to messages that were delayed because of SQS configurations like delay queues or retries. It doesn't reflect the immediate workload waiting to be processed.

upvoted 2 times

  **be9z** 12 months ago

The answer is D: This approach involves custom monitoring of queue depth and instance count, allowing dynamic scaling based on workload demands. By combining both metrics, you can achieve responsiveness while considering varying processing times.

upvoted 1 times

  **auxwww** 12 months ago

Selected Answer: D

ApproximateNumberOfMessagesDelayed - The number of messages in the queue that are delayed and not available for reading immediately. This can happen when the queue is configured as a delay queue or when a message has been sent with a delay parameter.

In our scenario the messages are not delayed in the queue - just not been picked up yet.

upvoted 1 times

  **JoeBoom** 1 year, 4 months ago

ApproximateNumberOfMessagesDelayed is a built in metric not a custom metric

Answer is B

upvoted 1 times

  **AminTriton** 3 months, 2 weeks ago

I guess you didn't read answer D properly. It says NumberOfMessagesVisible which is a different metric. So you can't answer based on custom vs built-in.

upvoted 1 times

🗨️ **Maria2023** 1 year, 4 months ago

Selected Answer: D

ApproximateNumberOfMessagesDelayed - The number of messages in the queue that are delayed and not available for reading immediately. This can happen when the queue is configured as a delay queue or when a message has been sent with a delay parameter.

In our scenario the messages are not delayed in the queue - just not been picked up yet. Hence I vote for "D"

upvoted 1 times

🗨️ **callspace** 1 year, 9 months ago

Selected Answer: D

Looks like the answer is in this link: (D)

ApproximateNumberOfMessagesDelayed:

The number of messages in the queue that are delayed and not available for reading immediately. This can happen when the queue is configured as a delay queue or when a message has been sent with a delay parameter. Which means not that they are delayed because of the insufficient available capacity to process but "delayed" is more of a feature rather than a result of something.

ApproximateNumberOfMessagesVisible: The number of messages to be processed.

upvoted 1 times

🗨️ **gokalpkoer3** 1 year, 10 months ago

Modifying an application is almost always a no-no in AWS exams. so I will go with B.

upvoted 2 times

🗨️ **jipark** 1 year, 10 months ago

Selected Answer: D

I thought B, but it looks D.

numberOfDelay seems just delay setting msg, not queued msg :

B. ApproximateNumberOfMessagesDelayed – Returns the approximate number of messages in the queue that are delayed and not available for reading immediately. This can happen when the queue is configured as a delay queue or when a message has been sent with a delay parameter.

D. ApproximateNumberOfMessagesNotVisible – Returns the approximate number of messages that are in flight. Messages are considered to be in flight if they have been sent to a client but have not yet been deleted or have not yet reached the end of their visibility window.

upvoted 4 times

🗨️ **hexie** 2 years ago

Selected Answer: B

if ApproximateNumberOfMessagesDelayed is increasing, wouldn't it be a good way to make the ASG scaling? Also providing the scalability accordingly to the problem it's facing? And "modifying the application to calculate.." doesn't also sound for something to the devs, not for the SysOps Administrator? I'm going for B.

upvoted 1 times

🗨️ **Vivec** 2 years, 3 months ago

Selected Answer: D

When there are delays in processing files due to a high volume of messages in the queue, adding more instances using Auto Scaling can help to reduce the processing time. The ApproximateNumberOfMessagesVisible metric is a good indicator of the workload on the EC2 instances. By creating an Auto Scaling policy based on this metric, the number of instances can be scaled up or down depending on the number of messages in the queue.

upvoted 2 times

🗨️ **Spike2020** 2 years, 4 months ago

C.

D: There is no "ApproximateNumberOfMessagesVisible" parameter. It should be ApproximateNumberOfMessagesNotVisible

upvoted 2 times

🗨️ **defmania00** 2 years, 4 months ago

There sure is an "ApproximateNumberOfMessagesVisible" parameter.

upvoted 2 times

🗨️ **jas26says** 1 year, 11 months ago

there isn't

upvoted 1 times

🗨️ 👤 **tttfakil** 1 year, 10 months ago

Why Not? <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-available-cloudwatch-metrics.html>

upvoted 2 times

🗨️ 👤 **squeeze_talus0y** 2 years, 5 months ago

Selected Answer: D

AWS offers even an example for this - <https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-target-tracking-metric-math.html#metric-math-sqs-queue-backlog>

upvoted 3 times

🗨️ 👤 **Mila28** 2 years, 6 months ago

Selected Answer: D

I'm agree with D

upvoted 1 times

🗨️ 👤 **foreverlearner** 2 years, 6 months ago

Selected Answer: D

The issue here is that the app is taking too long to process the files. This means that the messages are in the SQS queue, just they're not being processed fast enough.

It can't be B as "ApproximateNumberOfMessagesDelayed – Returns the approximate number of messages in the queue that are delayed and not available for reading immediately" (https://docs.aws.amazon.com/AWSSimpleQueueService/latest/APIReference/API_GetQueueAttributes.html).

This would mean that the messages are not available, which isn't the case in this question

The link posted by zolthar_z clearly explains why it's D. The app takes a variable amount of time to process each message, hence the app should calculate the Backlog per instance

upvoted 4 times

🗨️ 👤 **zolthar_z** 2 years, 6 months ago

Selected Answer: D

The answer is D: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-using-sqs-queue.html>

upvoted 3 times

🗨️ 👤 **marcelodba** 2 years, 7 months ago

Selected Answer: D

It's D

upvoted 1 times

🗨️ 👤 **marcelodba** 2 years, 6 months ago

Ans is B

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

upvoted 4 times

A company runs a multi-tier web application with two Amazon EC2 instances in one Availability Zone in the us-east-1 Region. A SysOps administrator must migrate one of the EC2 instances to a new Availability Zone. Which solution will accomplish this?

- A. Copy the EC2 instance to a different Availability Zone. Terminate the original instance.
- B. Create an Amazon Machine Image (AMI) from the EC2 instance and launch it in a different Availability Zone. Terminate the original instance.
- C. Move the EC2 instance to a different Availability Zone using the AWS CLI.
- D. Stop the EC2 instance, modify the Availability Zone, and start the instance.

Suggested Answer: B

Community vote distribution

B (100%)

  **princajen** Highly Voted 1 year, 4 months ago

Selected Answer: B

Answer is B!

It's not possible to move an existing instance to another subnet, Availability Zone, or VPC. Instead, you can manually migrate the instance by creating a new Amazon Machine Image (AMI) from the source instance. Then, launch a new instance using the new AMI in the desired subnet, Availability Zone, or VPC. Finally, you can reassign any Elastic IP addresses from the source instance to the new instance.

<https://aws.amazon.com/premiumsupport/knowledge-center/move-ec2-instance/>
upvoted 6 times

  **englishborn** Most Recent 8 months, 3 weeks ago



Selected Answer: B

B due to you cannot move an EC2's AV or VPC
upvoted 2 times

  **michaldavid** 1 year ago

Selected Answer: B

bbbbbbb
upvoted 1 times

  **Liongeek** 1 year, 1 month ago

Selected Answer: B

Ans: B
upvoted 2 times

A company is expanding its fleet of Amazon EC2 instances before an expected increase of traffic. When a SysOps administrator attempts to add more instances, an InstanceLimitExceeded error is returned.

What should the SysOps administrator do to resolve this error?

- A. Add an additional CIDR block to the VPC.
- B. Launch the EC2 instances in a different Availability Zone.
- C. Launch new EC2 instances in another VPC.
- D. Use Service Quotas to request an EC2 quota increase.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Hisayuki** Highly Voted 2 years, 4 months ago

Selected Answer: D

- Instance Limit Exceeded -> Over 64 vCPU
 - Insufficient Instance Capacity -> AWS's problem
 - Instance Terminates Immediately -> EBS volume limit, corrupt or no permission for decryption
 upvoted 21 times

🗳️ 👤 **princajen** Highly Voted 2 years, 10 months ago

Selected Answer: D

Description

You get the InstanceLimitExceeded error when you try to launch a new instance or restart a stopped instance.

Cause

If you get an InstanceLimitExceeded error when you try to launch a new instance or restart a stopped instance, you have reached the limit on the number of instances that you can launch in a Region. When you create your AWS account, we set default limits on the number of instances you can run on a per-Region basis.

Solution

You can request an instance limit increase on a per-region basis.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/troubleshooting-launch.html#troubleshooting-launch-limit>

upvoted 6 times

🗳️ 👤 **phuc12** Most Recent 5 months, 2 weeks ago

Selected Answer: D

Instance Limit Exceeded -> You have reached max vCPU limit in a Region
 C is wrong because a different VPC can still be in the same Region, and solving this problem requires you to launch EC2 in another Region
 upvoted 1 times

🗳️ 👤 **mk1523** 6 months ago

Selected Answer: D

Use Service Quotas to request an EC2 quota increase.
 upvoted 1 times

🗳️ 👤 **Reversing3314** 8 months, 3 weeks ago

If you get an InstanceLimitExceeded error when you try to launch a new instance or restart a stopped instance, you have reached the limit on the number of instances that you can launch in a Region. When you create your AWS account, we set default limits on the number of instances you can run on a per-Region basis.

Solution

You can request an instance limit increase on a per-region basis. For more information, see Amazon EC2 service quotas.


The option B gave AZ which is wrong it should be region
upvoted 1 times

  **MichalDavid** 2 years, 6 months ago

Selected Answer: D

ddddddd

upvoted 1 times

  **LionGeek** 2 years, 7 months ago

Ans: D

upvoted 1 times

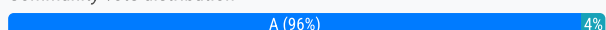
A company wants to prohibit its developers from using a particular family of Amazon EC2 instances. The company uses AWS Organizations and wants to apply the restriction across multiple accounts.

What is the MOST operationally efficient way for the company to apply service control policies (SCPs) to meet these requirements?

- A. Add the accounts to an organizational unit (OU). Apply the SCPs to the OU.
- B. Add the accounts to resource groups in AWS Resource Groups. Apply the SCPs to the resource groups.
- C. Apply the SCPs to each developer account
- D. Enroll the accounts with AWS Control Tower. Apply the SCPs to the AWS Control Tower management account.

Suggested Answer: A

Community vote distribution



Kinetix Highly Voted 2 years, 2 months ago

Selected Answer: A

Question says Most Operationally Efficient. If they do not already use CT it adds another layer of complexity. They will still need to create OU's for the Developers anyway.

upvoted 11 times

princajen Highly Voted 2 years, 4 months ago

Selected Answer: A

Answer is A!

<https://aws.amazon.com/blogs/industries/best-practices-for-aws-organizations-service-control-policies-in-a-multi-account-environment/>

upvoted 7 times

james2033 Most Recent 10 months, 1 week ago

Selected Answer: A

See <https://towardsthecloud.com/aws-scp-service-control-policies> . 'Attach' SCP (Services Control Policies) to OU (Organization Unit) is suitable word.

upvoted 2 times

jipark 1 year, 4 months ago

Selected Answer: A

why not B :

B."an" organizational unit is easier than A.resource "groups"

upvoted 1 times

michaldavid 2 years ago

Selected Answer: A

aaaaaaa

upvoted 2 times

Liongeek 2 years, 1 month ago

Selected Answer: A

Ans: A

Cause it's the most efficient

upvoted 3 times

bakjeeone 2 years, 2 months ago

Selected Answer: D

D, it is not required OU , it is sufficient by Control Tower

upvoted 1 times

Kinetix 2 years, 2 months ago

Question says Most Operationally Efficient. If they do not

already use CT it adds another layer of complexity. They will still need to create OU's for the Developers anyway. Its A

upvoted 7 times

An application is running on an Amazon EC2 instance in a VPC with the default DHCP option set. The application connects to an on-premises Microsoft SQL

Server database with the DNS name mssql.example.com. The application is unable to resolve the database DNS name.

Which solution will fix this problem?

- A. Create an Amazon Route 53 Resolver inbound endpoint. Add a forwarding rule for the domain example.com. Associate the forwarding rule with the VPC.
- B. Create an Amazon Route 53 Resolver inbound endpoint. Add a system rule for the domain example.com. Associate the system rule with the VPC.
- C. Create an Amazon Route 53 Resolver outbound endpoint. Add a forwarding rule for the domain example.com. Associate the forwarding rule with the VPC.
- D. Create an Amazon Route 53 Resolver outbound endpoint. Add a system rule for the domain example.com. Associate the system rule with the VPC.

Suggested Answer: C

Community vote distribution


C (100%)

 **Zulqarnain1** 6 months, 2 weeks ago

Selected Answer: C

Answer is C guys, have a look at this diagram <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html> and read the steps below. Outbound is when the AWS Infra is communicating to On-prem (via the resolver). Inbound is when on-prem is communicating with AWS Infra!

upvoted 1 times

 **Nazzhassan** 8 months, 2 weeks ago

A

Inbound endpoints are used to forward DNS queries from your network to Route 53 Resolver within your VPC. They specify the IP addresses that DNS resolvers on your network should forward DNS queries to. This is the opposite of what the scenario requires.

The scenario involves resolving DNS queries originating from an Amazon EC2 instance within a VPC. Therefore, the correct approach is to configure outbound forwarding, not inbound forwarding.

Option A correctly describes the solution by suggesting the creation of an outbound endpoint and adding a forwarding rule for the domain example.com to forward DNS queries from the VPC to DNS resolvers on the network.

upvoted 2 times

 **Zulqarnain1** 6 months, 2 weeks ago


Do you mean option C? Option A is referring to inbound and option C is referring to outbound

upvoted 1 times

 **Liongeek** 2 years, 1 month ago

Ans: C

upvoted 2 times

 **Atown** 2 years, 1 month ago

Selected Answer: C

C

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-forwarding-outbound-queries.html>

upvoted 4 times

 **princajen** 2 years, 3 months ago

Selected Answer: C

A!

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-rules-managing.html>

upvoted 1 times

  **princajen** 2 years, 3 months ago

Sorry C!

upvoted 5 times

  **kati2k22cz** 2 years, 3 months ago

Selected Answer: C

Checked its C.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-forwarding-outbound-queries.html>

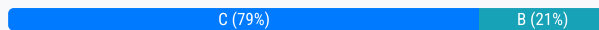
upvoted 2 times

A company's application is hosted by an internet provider at `app.example.com`. The company wants to access the application by using `www.company.com`, which the company owns and manages with Amazon Route 53. Which Route 53 record should be created to address this?

- A. A record
- B. Alias record
- C. CNAME record
- D. Pointer (PTR) record

Suggested Answer: C

Community vote distribution



🗳️ 👤 **OlehKom** Highly Voted 6 months, 2 weeks ago

Selected Answer: C

A record: Domain → IP address

Alias record: Domain → AWS resource

CNAME record: Domain → Another domain

PTR record: IP address → Domain (reverse lookup)

upvoted 9 times

🗳️ 👤 **kati2k22cz** Highly Voted 2 years, 3 months ago

C is correct.

"A CNAME record can redirect DNS queries to any DNS record. For example, you can create a CNAME record that redirects queries from `acme.example.com` to `zenith.example.com` or to `acme.example.org`."

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

upvoted 7 times

🗳️ 👤 **stoy123** Most Recent 10 months, 1 week ago

Selected Answer: B

B: Alias! It does the same as CNAME, it's cheaper and always preferred by AWS ;)

upvoted 1 times

🗳️ 👤 **AminTriton** 3 months, 2 weeks ago

It does not do the same as CNAME. Alias cannot point a hostname to another hostname.

upvoted 1 times

🗳️ 👤 **Zulqarnain1** 6 months, 2 weeks ago

Alias is for Domain to AWS Resource, not for Domain to Domain I believe.

upvoted 2 times

🗳️ 👤 **james2033** 10 months, 1 week ago

Selected Answer: C

It is CNAME for `www`, not just on AWS Route 53, for all DNS domain name servers in the world. See screenshot

<https://gist.github.com/assets/1328316/b8cc838c-3821-457e-ae0f-b3e0f70c8937>.

upvoted 1 times

🗳️ 👤 **jipark** 1 year, 4 months ago

Selected Answer: C

why not B :

- CNAME : hostname redirect to DNS

- Alias : hostname redirect to ARN

upvoted 2 times

🗳️ 👤 **hexie** 1 year, 6 months ago

Selected Answer: C

By what I understood, the website is hosted OUTSIDE AWS. Alias records aren't more suitable when mapping AWS resources to another? Like mapping an ALB to Route 53 alias.. Plus it's changing the root domain from example to company. I'm going for CNAME.

upvoted 1 times

🗨️ 👤 **braveheart22** 1 year, 10 months ago

C is the right way to go.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

upvoted 2 times

🗨️ 👤 **dasadanan** 1 year, 11 months ago

Selected Answer: C

I think the correct answer is c.

You should never use a CNAME record for your root domain name (e.g. example.com).

<https://support.dnsimple.com/articles/differences-between-a-cname-alias-url/>

upvoted 2 times

🗨️ 👤 **squeeze_talus0y** 1 year, 11 months ago

Selected Answer: B

If company.com is hosted in Route53 then Alias record can also be used.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html#resource-record-sets-choosing-alias-non-alias-comparison>

Alias records

An alias record can only redirect queries to selected AWS resources, such as the following:

Amazon S3 buckets

CloudFront distributions

Another record in the same Route 53 hosted zone

For example, you can create an alias record named acme.example.com that redirects queries to an Amazon S3 bucket that is also named acme.example.com. You can also create an acme.example.com alias record that redirects queries to a record named zenith.example.com in the example.com hosted zone.

upvoted 2 times

🗨️ 👤 **Spike2020** 1 year, 11 months ago

This website is hosted outside of AWS. It cannot be an alias.

upvoted 3 times

🗨️ 👤 **Gomer** 1 year, 7 months ago

app.example.com is hosted outside of AWS, but the company.com domain is hosted by Route 53. Why can't the www.company.com alias record point to any external domain? I'm not expert in DNS or Route 53, but it seems that if "company.com" is hosted by Route 53, the "www" record in Route 53 for that domain could be redirected (pointed) to anything with an alias (including cnn.com). I don't have to own a file to be able to create a link to it (even if I don't have rights to the file).

upvoted 1 times

🗨️ 👤 **BugsBunny9998666** 2 years ago

Selected Answer: C

URL to URL is CNAME. probably the only question in a whole sys ops where answer is CNAME

upvoted 4 times

🗨️ 👤 **michaldavid** 2 years ago

Selected Answer: C

cccccccc

upvoted 1 times

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

Ans: C

upvoted 1 times

A company expanded its web application to serve a worldwide audience. A SysOps administrator has implemented a multi-Region AWS deployment for all production infrastructure. The SysOps administrator must route traffic based on the location of resources. Which Amazon Route 53 routing policy should the SysOps administrator use to meet this requirement?

- A. Geolocation routing policy
- B. Geoproximity routing policy
- C. Latency-based routing policy
- D. Multivalue answer routing policy

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **haxaffee** Highly Voted 2 years, 9 months ago

Selected Answer: B

Should be B because it states "based on location of resources". Geolocation is used for based on location of users.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

upvoted 28 times

🗳️ 👤 **shure4shure** 2 years, 8 months ago

This is a great explanation

upvoted 4 times

🗳️ 👤 **jipark** 1 year, 10 months ago

I'll keep "geolocation for location of users"

upvoted 1 times

🗳️ 👤 **Zulqarnain1** Most Recent 6 months, 2 weeks ago

Selected Answer: B

Answer is B - Geoproximity routing policy as the question is asking based on the location of resources. NOT solely based on the location of users!

Location of Users = Geolocation

Location of Resources = Geoproximity

upvoted 3 times

🗳️ 👤 **Nazzhassan** 1 year, 2 months ago

C

In the context of a multi-region AWS deployment where the goal is to route traffic based on the location of resources, the most appropriate Amazon Route 53 routing policy is the Latency-based routing policy.

This policy evaluates the latency of the end user to different AWS regions and directs traffic to the region with the lowest latency for the user. By doing so, it ensures that users are routed to the region that provides the best performance for them, reducing latency and improving overall user experience.

Therefore, the correct option is:

C. Latency-based routing policy

upvoted 1 times

🗳️ 👤 **airraid2010** 2 years ago

Selected Answer: B

The Geolocation routing policy allows you to route traffic based on the geographic location of your users. You can create different routing configurations based on the continent, country, state, or even postal code of the requesting user. This policy is useful when you want to provide different responses or direct traffic to specific resources based on the user's location.

upvoted 1 times

🗨️ 👤 **Aamee** 8 months ago

You selected B for Geoproximity routing but explaining the justification of Geolocation routing policy instead lolz... :)
upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbb

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: B

Location of resources -based

upvoted 1 times

🗨️ 👤 **Atown** 2 years, 7 months ago

Selected Answer: B

Answer is B

Question "location of resources"

Geolocation routing policy – Use when you want to route traffic based on the location of your users. You can use geolocation routing to create records in a private hosted zone.

Geoproximity routing policy – Use when you want to route traffic based on the location of your resources and, optionally, shift traffic from resources in one location to resources in another.

upvoted 3 times

🗨️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: B

B. Geoproximity.

upvoted 1 times

🗨️ 👤 **princajen** 2 years, 9 months ago

Selected Answer: B

Geoproximity routing policy – Use when you want to route traffic based on the location of your resources and, optionally, shift traffic from resources in one location to resources in another.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

upvoted 1 times


A SysOps administrator wants to upload a file that is 1 TB in size from on-premises to an Amazon S3 bucket using multipart uploads. What should the SysOps administrator do to meet this requirement?

- A. Upload the file using the S3 console.
- B. Use the s3api copy-object command.
- C. Use the s3api put-object command.
- D. Use the s3 cp command

Suggested Answer: C

Community vote distribution

D (97%)

 **princajen** Highly Voted 2 years, 9 months ago

Selected Answer: D

I vote for D!

It's a best practice to use aws s3 commands (such as aws s3 cp) for multipart uploads and downloads, because these aws s3 commands automatically perform multipart uploading and downloading based on the file size. By comparison, aws s3api commands, such as aws s3api create-multipart-upload, should be used only when aws s3 commands don't support a specific upload need, such as when the multipart upload involves multiple servers, a multipart upload is manually stopped and resumed later, or when the aws s3 command doesn't support a required request parameter.


<https://aws.amazon.com/premiumsupport/knowledge-center/s3-multipart-upload-cli/>
upvoted 18 times

 **MrMLB** Highly Voted 2 years, 6 months ago

Selected Answer: D

The correct answer is D: Use the s3 cp command.

s3 cp file.txt s3://my-bucket/ --multipart-upload
upvoted 6 times

 **jipark** 1 year, 10 months ago
great explanation !!
upvoted 1 times

 **jagdishmav** Most Recent 7 months ago

Selected Answer: D

Important: It's a best practice to use aws s3 commands, such as aws s3 cp,

upvoted 1 times

 **Iamawstoo** 2 years, 5 months ago

A Sysops administrator wants to securely share an object from a private Amazon S3 bucket with a group of users who do not have an AWS account. What is the MOST operationally efficient solution that will meet this requirement?

- A. Attach an s3 bucket policy that only allows object downloads from the users' IP addresses.
 - B. Create an IAM role that has access to the object. Instruct the users to assume the role
 - C. Create an IAM user that has access to the object Share the credentials with the users
 - D. Generate a presigned URL for the object. Share the URL with the users.
- upvoted 1 times

 **Aamee** 9 months, 2 weeks ago

It is infact 'D' the correct ans. for this question
upvoted 1 times

🗳️ 👤 **XenonDemon** 1 year, 11 months ago

D is the answer

upvoted 2 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

dddddddd

upvoted 1 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Selected Answer: D

Ans: D

upvoted 1 times

🗳️ 👤 **Atown** 2 years, 7 months ago

Selected Answer: D

The Correct answer is D

Read Article

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-multipart-upload-cli/>

To use a high-level aws s3 command for your multipart upload, run this command:

```
$ aws s3 cp large_test_file s3://DOC-EXAMPLE-BUCKET/
```

Important: Use this aws s3api procedure only when aws s3 commands don't support a specific upload need

upvoted 1 times

🗳️ 👤 **[Removed]** 2 years, 9 months ago

I believe the correct answer is "D".

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/>

upvoted 1 times

🗳️ 👤 **Gorille69** 2 years, 9 months ago

It's a best practice to use aws s3 commands (such as aws s3 cp) for multipart uploads and downloads. here , there is a multipart uploads

upvoted 1 times

🗳️ 👤 **kati2k22cz** 2 years, 9 months ago

Selected Answer: C

Choose C. put-object is the answer.

<https://awscli.amazonaws.com/v2/documentation/api/latest/reference/s3api/put-object.html>

upvoted 1 times

🗳️ 👤 **haxaffee** 2 years, 9 months ago

Selected Answer: D

Ill go with D (aws s3 cp) on that one. It's recommended at <https://aws.amazon.com/premiumsupport/knowledge-center/s3-multipart-upload-cli/> and

there is more to be done for s3api to work.

upvoted 4 times

An application team is working with a SysOps administrator to define Amazon CloudWatch alarms for an application. The application team does not know the application's expected usage or expected growth. Which solution should the SysOps administrator recommend?

- A. Create CloudWatch alarms that are based on anomaly detection.
- B. Create CloudWatch alarms by using a set of composite alarms.
- C. Create CloudWatch alarms by using static thresholds.
- D. Create CloudWatch alarms that treat missing data as breaching.

Suggested Answer: A

Community vote distribution

A (100%)

🗲️ 👤 **kati2k22cz** Highly Voted 👍 9 months, 4 weeks ago

Selected Answer: A

Nice choice to use CloudWatch Anomaly Detection. Letter A

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Anomaly_Detection.html

upvoted 8 times

🗲️ 👤 **michaldavid** Most Recent ⌚ 6 months, 3 weeks ago

Selected Answer: A

aaaaaa

upvoted 1 times

🗲️ 👤 **Liongeek** 7 months, 2 weeks ago

Ans: A

upvoted 1 times

🗲️ 👤 **princajen** 9 months, 4 weeks ago

Selected Answer: A

A!

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Anomaly_Detection.html

upvoted 1 times

A company runs a stateless application that is hosted on an Amazon EC2 instance. Users are reporting performance issues. A SysOps administrator reviews the Amazon CloudWatch metrics for the application and notices that the instance's CPU utilization frequently reaches 90% during business hours. What is the MOST operationally efficient solution that will improve the application's responsiveness?

- A. Configure CloudWatch logging on the EC2 instance. Configure a CloudWatch alarm for CPU utilization to alert the SysOps administrator when CPU utilization goes above 90%.
- B. Configure an AWS Client VPN connection to allow the application users to connect directly to the EC2 instance private IP address to reduce latency.
- C. Create an Auto Scaling group, and assign it to an Application Load Balancer. Configure a target tracking scaling policy that is based on the average CPU utilization of the Auto Scaling group.
- D. Create a CloudWatch alarm that activates when the EC2 instance's CPU utilization goes above 80%. Configure the alarm to invoke an AWS Lambda function that vertically scales the instance.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ **Reversing3314** 8 months, 3 weeks ago
when question already states 90% utilization why choose average cpu percentage?
upvoted 1 times

🗳️ **michaldavid** 2 years, 6 months ago
Selected Answer: C
ccccccc
upvoted 2 times

🗳️ **Liongeek** 2 years, 7 months ago
Ans: C
upvoted 1 times

🗳️ **Atown** 2 years, 7 months ago
Selected Answer: C
I vote C
upvoted 3 times

🗳️ **kati2k22cz** 2 years, 9 months ago
Selected Answer: C
C
same example on this doc
<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-target-tracking.html>
upvoted 4 times

An ecommerce company uses an Amazon ElastiCache for Memcached cluster for in-memory caching of popular product queries on the shopping site. When viewing recent Amazon CloudWatch metrics data for the ElastiCache cluster, the SysOps administrator notices a large number of evictions.

Which of the following actions will reduce these evictions? (Choose two.)

- A. Add an additional node to the ElastiCache cluster.
- B. Increase the ElastiCache time to live (TTL).
- C. Increase the individual node size inside the ElastiCache cluster.
- D. Put an Elastic Load Balancer in front of the ElastiCache cluster.
- E. Use Amazon Simple Queue Service (Amazon SQS) to decouple the ElastiCache cluster.

Suggested Answer: BC

Community vote distribution

AC (98%)

haxaffee Highly Voted 2 years, 9 months ago

Selected Answer: AC

Scale Out and or Scale Up -> https://d1.awsstatic.com/training-and-certification/docs-sysops-associate/AWS-Certified-SysOps-Administrator-Associate_Sample-Questions_C02.pdf Question 2
upvoted 19 times

jipark Highly Voted 1 year, 10 months ago

Selected Answer: AC

what 'eviction' is :

Evictions occur when the cache is full and a new item needs to be added, resulting in the removal of an existing item from the cache.

upvoted 11 times

Reversing3314 Most Recent 8 months, 3 weeks ago

Why Not the Other Options?

B. Increase the ElastiCache time to live (TTL):

Increasing the TTL would actually keep data in the cache longer, potentially leading to more evictions since old data would stay in memory for longer periods, reducing available space for new data.

D. Put an Elastic Load Balancer in front of the ElastiCache cluster:

Elastic Load Balancers (ELBs) are used for distributing network traffic across EC2 instances, not for reducing cache evictions. They don't affect memory usage or cache management in ElastiCache.

E. Use Amazon SQS to decouple the ElastiCache cluster:

Amazon SQS is a message queue service and is unrelated to memory management or caching. It would not reduce evictions in ElastiCache

upvoted 2 times

tank7575 1 year, 4 months ago

Selected Answer: CE

chatGPT says C and E

upvoted 2 times

vivanchyk 1 year, 3 months ago

premium version? :)

the premium version says (and my brief humble research got matched with this):

"A. Add an additional node to the ElastiCache cluster.

Adding more nodes to the cluster increases the overall memory available for caching, reducing the likelihood of evictions. Each additional node

adds more memory capacity, allowing more data to be stored in the cache without evicting existing items.

C. Increase the individual node size inside the ElastiCache cluster.

By increasing the size of each node, you increase the amount of memory available in each node. This additional memory can accommodate more cached data, decreasing the need to evict older data when new data needs to be cached."

upvoted 1 times

🗨️ 👤 **AdamCzepiel** 1 year ago

meanwhile open source chaptgpt version also says it ^^

upvoted 2 times

🗨️ 👤 **Hatem08** 1 year, 7 months ago

Selected Answer: AC

AC scale up and scale out as the eviction is caused as the resources are not enough

upvoted 1 times

🗨️ 👤 **fig** 1 year, 7 months ago

Selected Answer: AC

The Evictions metric for Amazon ElastiCache for Memcached represents the number of non- expired items that the cache evicted to provide space for new items. If you are experiencing evictions with your cluster, it is usually a sign that you need to scale up (use a node that has a larger memory footprint) or scale out (add additional nodes to the cluster) to accommodate the additional data.

upvoted 2 times

🗨️ 👤 **tttfakil** 1 year, 10 months ago

Selected Answer: AC

AAAAACCCC

upvoted 2 times

🗨️ 👤 **bogossdu35** 2 years, 2 months ago

Selected Answer: AC

A and C

upvoted 2 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: AC

A and C

upvoted 2 times

🗨️ 👤 **Gianiluca** 2 years, 9 months ago

Selected Answer: AC

The is one of AWS sample exam questions and the answer is A C see - https://d1.awsstatic.com/training-and-certification/docs-sysops-associate/AWS-Certified-SysOps-Administrator-Associate_Sample-Questions.pdf

upvoted 6 times

A SysOps administrator wants to provide access to AWS services by attaching an IAM policy to multiple IAM users. The SysOps administrator also wants to be able to change the policy and create new versions.

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

- A. Add the users to an IAM service-linked role. Attach the policy to the role.
- B. Add the users to an IAM user group. Attach the policy to the group.
- C. Create an AWS managed policy.
- D. Create a customer managed policy.
- E. Create an inline policy.

Suggested Answer: BD

Community vote distribution

BD (100%)

  **kati2k22cz** Highly Voted 1 year, 3 months ago



Selected Answer: BD

Yes, B and D

<https://docs.aws.amazon.com/acm/latest/userguide/authen-custmanagedpolicies.html>

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups.html

upvoted 8 times

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

Selected Answer: BD

BD. Don't use AWS managed policies because you can't customise those.



upvoted 5 times

  **wooyourdaddy** Most Recent 8 months ago

Selected Answer: BD

Ref link: https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_managed-vs-inline.html#customer-managed-policies

upvoted 1 times

  **bogossdu35** 8 months, 3 weeks ago

Selected Answer: BD

I agree, B and D

upvoted 2 times

A company stores critical data in Amazon S3 buckets. A SysOps administrator must build a solution to record all S3 API activity. Which action will meet this requirement?

- A. Configure S3 bucket metrics to record object access logs.
- B. Create an AWS CloudTrail trail to log data events for all S3 objects.
- C. Enable S3 server access logging for each S3 bucket.
- D. Use AWS IAM Access Analyzer for Amazon S3 to store object access logs.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **princajen** Highly Voted 1 year, 9 months ago

Selected Answer: B

B!

Amazon S3 is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Amazon S3. CloudTrail captures a subset of API calls for Amazon S3 as events, including calls from the Amazon S3 console and code calls to the Amazon S3 APIs.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging.html>

upvoted 5 times

🗳️ 👤 **chanaka5** 1 year, 9 months ago

Question specifies API calls, hence S3 Server access logging is out.

upvoted 3 times

🗳️ 👤 **jipark** Most Recent 10 months, 3 weeks ago

Selected Answer: B

not A. S3 bucket metrics do not record API

not C. server access logging include only access log

not D. IAM Access Analyzer store only IAM Access log

upvoted 2 times

🗳️ 👤 **s50600822** 1 year ago

bad question, there should be a "both", looking at the table on <https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html>

one can argue what are more important but it's very fuzzy

upvoted 1 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

I'll go B

upvoted 2 times

🗳️ 👤 **Surferbolt** 1 year, 8 months ago

Selected Answer: B

B. CloudTrail logs API calls

upvoted 2 times

A company runs an application that uses a MySQL database on an Amazon EC2 instance. The EC2 instance has a General Purpose SSD Amazon Elastic Block

Store (Amazon EBS) volume. The company made changes to the application code and now wants to perform load testing to evaluate the impact of the code changes.

A SysOps administrator must create a new MySQL instance from a snapshot of the existing production instance. This new instance needs to perform as similarly as possible to the production instance.

Which restore option meets these requirements?

- A. Use EBS fast snapshot restore to create a new General Purpose SSD EBS volume from the production snapshot.
- B. Use EBS fast snapshot restore to create a new Provisioned IOPS SSD EBS volume from the production snapshot.
- C. Use EBS snapshot restore to create a new General Purpose SSD EBS volume from the production snapshot.
- D. Use EBS snapshot restore to create a new Provisioned IOPS SSD EBS volume from the production snapshot.

Suggested Answer: A

Community vote distribution

A (80%)

C (20%)

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

Selected Answer: A


Seems to me that A is for exam only.

For real life it is C .

There are no requirements on how fast EBS should provide its base performance.

If QA team will start using new instance tomorrow - then A is unneeded overhead.

upvoted 14 times

 **AminTriton** Most Recent 3 months, 2 weeks ago

Selected Answer: A

If you have answered C, remember that if the restore takes a long time for a production environment, then you might lose some data that has been modified or added to your original database during this long period of restore. Therefore, it is best to do it quickly and use FSR although FSR would cost slightly more.

upvoted 1 times

 **numark** 7 months, 1 week ago

A>>> FSR allows for faster recovery times when restoring EBS snapshots, which is important when you're performing load testing and need the new environment set up quickly. FSR enables the ability to restore data from an Amazon EBS snapshot and have it be immediately available for use with minimal delay, providing you with better performance during the recovery process.

GP3 Volume: restoring the new MySQL instance on a gp3 volume would ensure the performance characteristics are as similar as possible to the production setup. This gives you a more accurate environment for load testing the impact of the code changes. gp3 volumes are suitable for most use cases where the performance requirements are not excessively high.

C>>> While this option will restore a gp3 volume from the snapshot, it does not include Fast Snapshot Restore. Without FSR, the restoration process might take longer, which is less ideal for quickly setting up a new test environment.

upvoted 2 times

 **VerRi** 11 months, 2 weeks ago

Selected Answer: A

A, After restoring the volume from the snapshot, data remains in S3 until the first read, causing initial performance degradation. Fast Snapshot Restore will preload the data upon restoration.

upvoted 2 times

 **Maria2023** 1 year, 4 months ago

Selected Answer: C

Athough we all

upvoted 1 times

 **johnnyjin** 1 year, 9 months ago

C, only Volumes provisioned with performance up to 64,000 IOPS and 1,000 MiB/s throughput receive the full performance benefit of fast snapshot restore.

upvoted 1 times

🗨️ 👤 **xSohox** 1 year, 10 months ago

Selected Answer: A

Definitely A:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/backup-recovery/restore.html#restore-snapshot>

"If an application accesses the volume where the data is not loaded, there is higher latency than normal while the data is loaded from Amazon S3. To avoid this impact for latency-sensitive applications, you have two options:

-You can initialize the EBS volume.

-For an additional charge, Amazon EBS supports fast snapshot restore, which eliminates the need initialize your volume."

upvoted 4 times

🗨️ 👤 **jipark** 1 year, 10 months ago

Selected Answer: C

A or C :

there's no clue about fast or normal snapshot.

upvoted 4 times

🗨️ 👤 **s50600822** 2 years ago

it's weird... one can argue we need fast snapshot so the latency between prod and this is minimal but ...

upvoted 3 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaa

upvoted 1 times

🗨️ 👤 **Liongeek** 2 years, 7 months ago

Ans: A

upvoted 1 times

🗨️ 👤 **princajen** 2 years, 10 months ago

Selected Answer: A

The answer is A!

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-fast-snapshot-restore.html>

upvoted 4 times

A team of on-call engineers frequently needs to connect to Amazon EC2 instances in a private subnet to troubleshoot and run commands. The instances use either the latest AWS-provided Windows Amazon Machine Images (AMIs) or Amazon Linux AMIs. The team has an existing 1AM role for authorization. A SysOps administrator must provide the team with access to the instances by granting IAM permissions to this role.

Which solution will meet this requirement?

- A. Add a statement to the 1AM role policy to allow the `ssm:StartSession` action on the instances. Instruct the team to use AWS Systems Manager Session Manager to connect to the instances by using the assumed IAM role.
- B. Associate an Elastic IP address and a security group with each instance. Add the engineers' IP addresses to the security group inbound rules. Add a statement to the IAM role policy to allow the `ec2:AuthorizeSecurityGroupIngress` action so that the team can connect to the instances.
- C. Create a bastion host with an EC2 instance, and associate the bastion host with the VPC. Add a statement to the 1AM role policy to allow the `ec2:CreateVpnConnection` action on the bastion host. Instruct the team to use the bastion host endpoint to connect to the instances.
- D. Create an internet-facing Network Load Balancer. Use two listeners. Forward port 22 to a target group of Linux instances. Forward port 3389 to a target group of Windows instances. Add a statement to the IAM role policy to allow the `ec2:CreateRoute` action so that the team can connect to the instances.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaaa

upvoted 3 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

Ans: A

upvoted 1 times

🗳️ 👤 **Surferbolt** 1 year, 8 months ago

Selected Answer: A

A is correct.

upvoted 2 times

🗳️ 👤 **kati2k22cz** 1 year, 9 months ago

Selected Answer: A

Agree, its A.

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

upvoted 2 times

🗳️ 👤 **jipark** 10 months, 3 weeks ago

AWS Systems Manager Session Manager: This service allows you to manage interactive sessions to EC2 instances without needing to open inbound ports, manage bastion hosts, or assign Elastic IPs. It provides secure and controlled access to instances.

upvoted 2 times

A company needs to ensure strict adherence to a budget for 25 applications deployed on AWS. Separate teams are responsible for storage, compute, and database costs. A SysOps administrator must implement an automated solution to alert each team when their projected spend will exceed a quarterly amount that has been set by the finance department. The solution cannot incur additional compute, storage, or database costs. Which solution will meet these requirements?

- A. Configure AWS Cost and Usage Reports to send a daily report to an Amazon S3 bucket. Create an AWS Lambda function that will evaluate spend by service and notify each team by using Amazon Simple Notification Service (Amazon SNS) notifications. Invoke the Lambda function when a report is placed in the S3 bucket.
- B. Configure AWS Cost and Usage Reports to send a daily report to an Amazon S3 bucket. Create a rule in Amazon EventBridge (Amazon CloudWatch Events) to evaluate the spend by service and notify each team by using Amazon Simple Queue Service (Amazon SQS) when the cost threshold is exceeded.
- C. Use AWS Budgets to create one cost budget and select each of the services in use. Specify the budget amount defined by the finance department along with the forecasted cost threshold. Enter the appropriate email recipients for the budget.
- D. Use AWS Budgets to create a cost budget for each team, filtering by the services they own. Specify the budget amount defined by the finance department along with a forecasted cost threshold. Enter the appropriate email recipients for each budget.

Suggested Answer: D

Community vote distribution

D (100%)

fig 7 months, 2 weeks ago

Selected Answer: D

Don't like this!

Budgets do not guarantee to stop you from exceeding a spend. But D seems the best option from those presented.

There can be a delay between when you incur a charge and when you receive a notification from AWS Budgets for the charge. This is due to a delay between when an AWS resource is used and when that resource usage is billed. You might incur additional costs or usage that exceed your budget notification threshold before AWS Budgets can notify you.

<https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-managing-costs.html>

upvoted 1 times

englishborn 1 year, 2 months ago

Selected Answer: D

C is services in use vs D which is services each team own

upvoted 3 times

jipark 10 months, 3 weeks ago

I agree

upvoted 1 times

michaldavid 1 year, 6 months ago

Selected Answer: D

dddddd

upvoted 2 times

Liongeek 1 year, 7 months ago

Ans: D

You will need one budget per department

upvoted 3 times

Surferbolt 1 year, 8 months ago

Selected Answer: D

D. Budgets

upvoted 2 times

🗨️ 👤 **kati2k22cz** 1 year, 9 months ago

Selected Answer: D

Letter D. Cost per team

<https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-managing-costs.html>

upvoted 4 times

A company hosts a static website on Amazon S3. An Amazon CloudFront distribution presents this site to global users. The company uses the Managed-CachingDisabled CloudFront cache policy. The company's developers confirm that they frequently update a file in Amazon S3 with new information.

Users report that the website presents correct information when the website first loads the file. However, the users' browsers do not retrieve the updated file after a refresh.

What should a SysOps administrator recommend to fix this issue?

- A. Add a Cache-Control header field with max-age=0 to the S3 object.
- B. Change the CloudFront cache policy to Managed-CachingOptimized.
- C. Disable bucket versioning in the S3 bucket configuration.
- D. Enable content compression in the CloudFront configuration.

Suggested Answer: A

Community vote distribution

A (93%)

7%

 **princajen** Highly Voted 2 years, 9 months ago

Selected Answer: A

A!

You can control how long your files stay in a CloudFront cache before CloudFront forwards another request to your origin. Reducing the duration allows you to serve dynamic content. Increasing the duration means that your users get better performance because your files are more likely to be served directly from the edge cache. A longer duration also reduces the load on your origin.

To change the cache duration for an individual file, you can configure your origin to add a Cache-Control header with the max-age or s-maxage directive, or an Expires header to the file.

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

upvoted 6 times

 **be9z** Most Recent 12 months ago

By adding a Cache-Control header with max-age=0 to the S3 object, you instruct CloudFront to revalidate the content on every request.

When a user refreshes the page, CloudFront checks with S3 to ensure it has the latest version of the file.

This approach ensures that users receive the most up-to-date content.


upvoted 2 times

 **bruppp31** 1 year, 7 months ago

I think this question is wrong. By having the cache policy set to Managed-CachingDisabled all files will be fetched every time. This is also true of A if we have Managed-CachingOptimized.

The answer is probably A but it shouldn't make a difference in this case.

upvoted 2 times

 **callspace** 1 year, 9 months ago

Selected Answer: A

link in the comments section states:

The Cache-Control max-age directive lets you specify how long (in seconds) that you want an object to remain in the cache before CloudFront gets the object again from the origin server. The minimum expiration time CloudFront supports is 0 seconds. The maximum value is 100 years. And the company continue using the Managed-CachingDisabled CloudFront cache policy.

upvoted 3 times

 **Hisayuki** 2 years, 3 months ago

Selected Answer: A

You should disable Cache in Browser with Cache-Control: max-age=0.



upvoted 2 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaaaaa

upvoted 1 times

  **Atown** 2 years, 7 months ago

Selected Answer: B

Would it not be B since Caching is already Disabled?

upvoted 1 times

  **Surferbolt** 2 years, 8 months ago

Selected Answer: A

A is the answer.

upvoted 2 times

A company has a policy that requires all Amazon EC2 instances to have a specific set of tags. If an EC2 instance does not have the required tags, the noncompliant instance should be terminated.

What is the MOST operationally efficient solution that meets these requirement?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to send all EC2 instance state changes to an AWS Lambda function to determine if each instance is compliant. Terminate any noncompliant instances.
- B. Create an IAM policy that enforces all EC2 instance tag requirements. If the required tags are not in place for an instance, the policy will terminate noncompliant instance.
- C. Create an AWS Lambda function to determine if each EC2 instance is compliant and terminate an instance if it is noncompliant. Schedule the Lambda function to invoke every 5 minutes.
- D. Create an AWS Config rule to check if the required tags are present. If an EC2 instance is noncompliant, invoke an AWS Systems Manager Automation document to terminate the instance.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **johnson_chao** 9 months ago

Selected Answer: D

DDDDDD

upvoted 1 times

🗳️ 👤 **ronnykapo** 1 year, 6 months ago

DDDDDD

upvoted 2 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: D

dddddd

upvoted 2 times

🗳️ 👤 **Liongeek** 2 years, 1 month ago

Ans: D

upvoted 1 times

🗳️ 👤 **Surferbolt** 2 years, 2 months ago

Selected Answer: D

D. Use Config.

upvoted 1 times

🗳️ 👤 **kati2k22cz** 2 years, 3 months ago

its D.

https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_use-managed-rules.html

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

upvoted 4 times

🗳️ 👤 **jipark** 1 year, 4 months ago

Config Rule terminate or restrict many rules.

upvoted 1 times

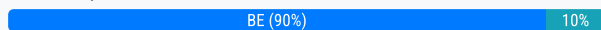
A SysOps administrator wants to manage a web server application with AWS Elastic Beanstalk. The Elastic Beanstalk service must maintain full capacity for new deployments at all times.

Which deployment policies satisfy this requirement? (Choose two.)

- A. All at once
- B. Immutable
- C. Rebuild
- D. Rolling
- E. Rolling with additional batch

Suggested Answer: BE

Community vote distribution



Atown Highly Voted 2 years, 1 month ago

Selected Answer: BE

BE

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

B -- Immutable deployments perform an immutable update to launch a full set of new instances running the new version of the application in a separate Auto Scaling group, alongside the instances running the old version. Immutable deployments can prevent issues caused by partially completed rolling deployments. If the new instances don't pass health checks, Elastic Beanstalk terminates them, leaving the original instances untouched.

E --To maintain full capacity during deployments, you can configure your environment to launch a new batch of instances before taking any instances out of service. This option is known as a rolling deployment with an additional batch. When the deployment completes, Elastic Beanstalk terminates the additional batch of instances.

upvoted 12 times

princajen Highly Voted 2 years, 4 months ago

Selected Answer: BE

Vote for B,E

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

upvoted 5 times

konieczny69 Most Recent 1 year ago

Selected Answer: CD

Not in the scope of the exam

upvoted 3 times

michaldavid 2 years ago

Selected Answer: BE

B and E

upvoted 2 times

Liongeek 2 years, 1 month ago

Ans: BE

upvoted 1 times

A company has an Auto Scaling group of Amazon EC2 instances that scale based on average CPU utilization. The Auto Scaling group events log indicates an `InsufficientInstanceCapacity` error.

Which actions should a SysOps administrator take to remediate this issue? (Choose two.)

- A. Change the instance type that the company is using.
- B. Configure the Auto Scaling group in different Availability Zones.
- C. Configure the Auto Scaling group to use different Amazon Elastic Block Store (Amazon EBS) volume sizes.
- D. Increase the maximum size of the Auto Scaling group.
- E. Request an increase in the instance service quota.

Suggested Answer: BE

Community vote distribution

AB (100%)


  **haxaffee** Highly Voted 1 year, 9 months ago

Selected Answer: AB

Its AB. E is relevant for `InstanceLimitExceeded` not `InsufficientInstanceCapacity`.
upvoted 22 times

  **jipark** 10 months, 3 weeks ago

thanks ton of it !!
upvoted 2 times

  **skinnyxuppy** Highly Voted 1 year, 10 months ago

Selected Answer: AB

<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/troubleshooting-launch.html#troubleshooting-launch-capacity>

Solution

To resolve the issue, try the following:

Wait a few minutes and then submit your request again; capacity can shift frequently.



Submit a new request with a reduced number of instances. For example, if you're making a single request to launch 15 instances, try making 3 requests for 5 instances, or 15 requests for 1 instance instead.

****If you're launching an instance, submit a new request without specifying an Availability Zone.

****If you're launching an instance, submit a new request using a different instance type (which you can resize at a later stage). For more information, see [Change the instance type](#).

If you are launching instances into a cluster placement group, you can get an insufficient capacity error. For more information, see [Placement group rules and limitations](#).

upvoted 8 times

  **Christina666** Most Recent 11 months, 1 week ago

Selected Answer: AB

Cause

If you get this error when you try to launch an instance or restart a stopped instance, AWS does not currently have enough available On-Demand capacity to fulfill your request.

upvoted 2 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: AB

A and B

upvoted 2 times

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

Ans: B

upvoted 1 times

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

Selected Answer: AB

AB is the answer.

upvoted 2 times

🗨️ 👤 **princajen** 1 year, 9 months ago

Selected Answer: AB

It's A and B. They had another question asking the same thing and the answers were A and B.

upvoted 2 times

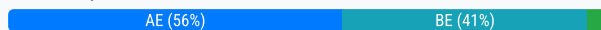
A SysOps administrator needs to control access to groups of Amazon EC2 instances using AWS Systems Manager Session Manager. Specific tags on the EC2 instances have already been added.

Which additional actions should the administrator take to control access? (Choose two.)

- A. Attach an IAM policy to the users or groups that require access to the EC2 instances.
- B. Attach an IAM role to control access to the EC2 instances.
- C. Create a placement group for the EC2 instances and add a specific tag.
- D. Create a service account and attach it to the EC2 instances that need to be controlled.
- E. Create an IAM policy that grants access to any EC2 instances with a tag specified in the Condition element.

Suggested Answer: BE

Community vote distribution



Vivec Highly Voted 1 year, 9 months ago

Selected Answer: AE

A. Attach an IAM policy to the users or groups that require access to the EC2 instances: IAM policies can be used to control access to resources in AWS. The policy can specify which actions are allowed or denied and which resources the user or group can access. In this case, the policy should include permissions to use the Session Manager service.

E. Create an IAM policy that grants access to any EC2 instances with a tag specified in the Condition element: This policy can specify that access is granted only to instances with specific tags. For example, a policy could specify that users or groups can only access instances that have a specific tag, such as "Environment=Prod". This helps to ensure that only the appropriate instances are accessed.

upvoted 7 times

OlehKom Most Recent 6 months, 2 weeks ago

Selected Answer: AE

Simply remember: The IAM role attached to ec2 instances (B) allows the instances to interact with AWS services (Systems Manager) but doesn't control user access

upvoted 2 times

AgboolaKun 7 months, 1 week ago

Selected Answer: AE

Control access to smaller deployments of Amazon EC2 instances as follows:

1. Add a specific tag to the instances that you want to grant the users or groups access to.
2. Create an IAM policy that grants access to any instances with the specific tag.
3. Attach the IAM policy to the users or groups that you want to access the instances.

Please refer to this link for more information - <https://repost.aws/knowledge-center/iam-ec2-resource-tags>

upvoted 2 times

stoy123 10 months, 1 week ago

Selected Answer: BE

B and E

upvoted 2 times

Maria2023 10 months, 3 weeks ago

Selected Answer: BE

The question does not mention granting users access to anything. Here we only discuss the access from SSM to the instances, which is not enabled by default.

upvoted 2 times

Rabbit117 11 months, 2 weeks ago

Selected Answer: AE

You would create an IAM policy that grants access to any EC2 instance with a tag specified in the condition element (E) and then you would attach that policy to the users or groups which require the access (A). I think B is wrong as you assume an IAM role, you don't attach it...

upvoted 2 times

🗳️ 👤 **xdkonorek2** 11 months, 3 weeks ago

Selected Answer: BE

To achieve this for all the instances without considering tags you would create an instance profile with a managed policy - "AmazonSSMMManagedInstanceCore" and attach it to an instance.

For instances with specified tags you must create your own conditional policy so SSM agent has access only to instances with particular tags.

upvoted 1 times

🗳️ 👤 **konieczny69** 1 year ago

Selected Answer: AE

A and E

upvoted 3 times

🗳️ 👤 **Hatem08** 1 year, 1 month ago

Selected Answer: AE

AE is the correct answer !

upvoted 3 times

🗳️ 👤 **wh1t4k3r** 1 year, 3 months ago

I go for A and E, given that creating a policy and attaching a role does not solve the problem. You need to attach the policy you've created somewhere.

upvoted 3 times

🗳️ 👤 **xSohox** 1 year, 4 months ago

Selected Answer: BE

I think it is BE.

Because "By default, AWS Systems Manager doesn't have permission to perform actions on your instances."

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

So you need to grant Session Manager the permission to perform actions on your Amazon EC2 instances:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/getting-started-create-iam-instance-profile.html>

upvoted 1 times

🗳️ 👤 **konieczny69** 1 year ago

it can't be BE since they are logically unrelated, where one creates a role and second a policy

it's AE

upvoted 2 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: BE

A & E is a pair, can't be both

upvoted 1 times

🗳️ 👤 **konieczny69** 1 year ago

it must be both, create and attach are inseparable

AE

upvoted 2 times

🗳️ 👤 **Gomer** 1 year, 6 months ago

Selected Answer: BE

Could have been A, except you wouldn't attach a policy to a "users or a groups". That's too broad. Best practice is to use a group or a role, and not apply policies directly to a user.

upvoted 3 times

🗳️ 👤 **fazlur21** 1 year, 6 months ago

The question is "A SysOps Administrator needs to control access __to__ groups of Amazon EC2 instances.", not the control access of EC2 instances (or from them), so B is not applicable here.

upvoted 2 times

🗳️ 👤 **fazlur21** 1 year, 6 months ago

BE

is the real answer 100%

upvoted 1 times

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

Selected Answer: BE

```
"Condition": {  
  "StringEquals": {  
    "ec2:ResourceTag/TAG-KEY": "TAG-VALUE"  
  }  
}
```

<https://repost.aws/knowledge-center/iam-policy-permission-ec2-tags-vpc>
upvoted 2 times

🗨️ 👤 **Gomer** 1 year, 8 months ago

Selected Answer: BE

Q97: A team of on-call engineers frequently needs to connect to Amazon EC2 instances in a private subnet to troubleshoot and run commands. The instances use either the latest AWS-provided Windows Amazon Machine Images (AMIs) or Amazon Linux AMIs. The team has an existing 1AM role for authorization. A SysOps administrator must provide the team with access to the instances by granting IAM permissions to this role. Which solution will meet this requirement?"

Correct answer: A. Add a statement to the 1AM role policy to allow the ssm:StartSession action on the instances. Instruct the team to use AWS Systems Manager Session Manager to connect to the instances by using the assumed IAM role."

None of the other answers have anything to do with a policy for group or users.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/getting-started-restrict-access-examples.html#restrict-access-example-instance-tags>

I vote for "BE" because of policy example and I know role could work. I think wording in "A." is purposefully vague.

upvoted 3 times

A company has an AWS Lambda function in Account A. The Lambda function needs to read the objects in an Amazon S3 bucket in Account B. A SysOps administrator must create corresponding IAM roles in both accounts. Which solution will meet these requirements?

- A. In Account A, create a Lambda execution role to assume the role in Account B. In Account B, create a role that the function can assume to gain access to the S3 bucket.
- B. In Account A, create a Lambda execution role that provides access to the S3 bucket. In Account B, create a role that the function can assume.
- C. In Account A, create a role that the function can assume. In Account B, create a Lambda execution role that provides access to the S3 bucket.
- D. In Account A, create a role that the function can assume to gain access to the S3 bucket. In Account B, create a Lambda execution role to assume the role in Account A.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **kati2k22cz** Highly Voted 9 months, 3 weeks ago

Selected Answer: A

A.

<https://aws.amazon.com/premiumsupport/knowledge-center/lambda-function-assume-iam-role/>

upvoted 6 times

🗳️ **OlehKom** Most Recent 6 months, 2 weeks ago

Selected Answer: A

Account A: Lambda function needs a role to "ask permission" (assume a role) in Account B.

Account B: S3 bucket has a role that says, "If Account A asks properly, let them access my bucket."

upvoted 2 times

🗳️ **michaldavid** 6 months, 3 weeks ago

Selected Answer: A

aaaaaaaa

upvoted 3 times

🗳️ **Liongeek** 7 months, 2 weeks ago

Ans: A

upvoted 1 times

🗳️ **Surferbolt** 8 months, 2 weeks ago

Selected Answer: A

A is the answer.

upvoted 1 times

An AWS Lambda function is intermittently failing several times a day. A SysOps administrator must find out how often this error has occurred in the last 7 days.

Which action will meet this requirement in the MOST operationally efficient manner?

- A. Use Amazon Athena to query the Amazon CloudWatch logs that are associated with the Lambda function.
- B. Use Amazon Athena to query the AWS CloudTrail logs that are associated with the Lambda function.
- C. Use Amazon CloudWatch Logs Insights to query the associated Lambda function logs.
- D. Use Amazon OpenSearch Service (Amazon Elasticsearch Service) to stream the Amazon CloudWatch logs for the Lambda function.

Suggested Answer: A

Community vote distribution

C (100%)

🗳️ 👤 **jipark** Highly Voted 1 year, 10 months ago

Selected Answer: C

why not A :

Using Amazon Athena to query CloudWatch logs could work, but it introduces additional complexity by setting up and managing Athena.

why not B :

Querying AWS CloudTrail logs is primarily for tracking API activity and resource changes. It might not provide the same level of visibility into Lambda function errors

upvoted 8 times

🗳️ 👤 **princajen** Highly Voted 2 years, 9 months ago

Selected Answer: C

Voting for C!

<https://aws.amazon.com/blogs/compute/operating-lambda-using-cloudwatch-logs-insights/>

upvoted 6 times

🗳️ 👤 **auxwww** Most Recent 12 months ago

Selected Answer: C

Log insights is the easiest way to access this info. Athen doesn't query CW logs directly. Using this will require exporting CW logs, uploading to S3, creating a table which is least optimal solution. Therefore C

upvoted 3 times

🗳️ 👤 **stoy123** 1 year, 4 months ago

Selected Answer: C

C .

upvoted 1 times

🗳️ 👤 **hexie** 2 years ago

Selected Answer: C

Well it doesnt makes sense to be A since it would be needed to configure a Bucket for the logs plus setting up Athena for querying the Bucket. I'm going for C.

upvoted 3 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: C

cccccc

upvoted 1 times

🗳️ 👤 **Liongeek** 2 years, 7 months ago

Ans: C



upvoted 2 times

🗳️ 👤 **Surferbolt** 2 years, 8 months ago

Selected Answer: C

C is the answer

upvoted 3 times

  **Rick365** 2 years, 10 months ago

Selected Answer: C

C. Use Amazon CloudWatch Logs Insights to query the associated Lambda function logs.

upvoted 3 times

A company is using Amazon CloudFront to serve static content for its web application to its users. The CloudFront distribution uses an existing on-premises website as a custom origin. The company requires the use of TLS between CloudFront and the origin server. This configuration has worked as expected for several months. However, users are now experiencing HTTP 502 (Bad Gateway) errors when they view webpages that include content from the CloudFront distribution.

What should a SysOps administrator do to resolve this problem?

- A. Examine the expiration date on the certificate on the origin site. Validate that the certificate has not expired. Replace the certificate if necessary.
- B. Examine the hostname on the certificate on the origin site. Validate that the hostname matches one of the hostnames on the CloudFront distribution. Replace the certificate if necessary.
- C. Examine the firewall rules that are associated with the origin server. Validate that port 443 is open for inbound traffic from the internet. Create an inbound rule if necessary.
- D. Examine the network ACL rules that are associated with the CloudFront distribution. Validate that port 443 is open for outbound traffic to the origin server. Create an outbound rule if necessary.

Suggested Answer: A

Community vote distribution

A (100%)

  **princajen**  2 years, 3 months ago

Selected Answer: A

A! is correct

HTTP 502 errors from CloudFront can occur because of the following reasons:

There's an SSL negotiation failure because the origin is using SSL/TLS protocols and ciphers that aren't supported by CloudFront.
 There's an SSL negotiation failure because the SSL certificate on the origin is expired or invalid, or because the certificate chain is invalid.
 There's a host header mismatch in the SSL negotiation between your CloudFront distribution and the custom origin.
 The custom origin isn't responding on the ports specified in the origin settings of the CloudFront distribution.
 The custom origin is ending the connection to CloudFront too quickly.

<https://aws.amazon.com/premiumsupport/knowledge-center/resolve-cloudfront-connection-error/>

upvoted 16 times

  **tamng** 1 year ago

great !

upvoted 1 times

  **konieczny69**  1 year ago

Selected Answer: A

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/http-502-bad-gateway.html#origin-not-responding-on-specified-ports>

could be C, however the keyword is that it used to work and stopped
 hence A

upvoted 2 times

  **bruppp31** 1 year, 1 month ago

I think this could be A or B. Both are equally as likely. The certificate could either be expired or doesn't match.

Reference: <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/http-502-bad-gateway.html>

upvoted 1 times

  **Christina666** 1 year, 5 months ago

Selected Answer: A

An HTTP 502 status code (Bad Gateway) indicates that CloudFront wasn't able to serve the requested object because it couldn't connect to the origin server.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/http-502-bad-gateway.html>

upvoted 2 times

  **braveheart22** 1 year, 10 months ago

A is the way here

upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: A

aaaaaaaaa

upvoted 1 times

An Amazon CloudFront distribution has a single Amazon S3 bucket as its origin. A SysOps administrator must ensure that users can access the S3 bucket only through requests from the CloudFront endpoint.
Which solution will meet these requirements?

- A. Configure S3 Block Public Access on the S3 bucket. Update the S3 bucket policy to allow the GetObject action from only the CloudFront distribution.
- B. Configure Origin Shield in the CloudFront distribution. Update the CloudFront origin to include a custom Origin_Shield header.
- C. Create an origin access identity (OAI). Assign the OAI to the CloudFront distribution. Update the S3 bucket policy to restrict access to the OAI.
- D. Create an origin access identity (OAI). Assign the OAI to the S3 bucket. Update the CloudFront origin to include a custom Origin header with the OAI value.

Suggested Answer: C

Community vote distribution

C (100%)

  **kati2k22cz** Highly Voted 1 year, 9 months ago

Selected Answer: C

C is the correct answer

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>

upvoted 5 times

  **jipark** Most Recent 10 months, 3 weeks ago

Selected Answer: C

Origin Access Identity (OAI): An OAI is a special CloudFront user that you can associate with one or more CloudFront distributions



upvoted 2 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: C

ccccccc


upvoted 2 times

  **Surferbolt** 1 year, 8 months ago

Selected Answer: C

C is the answer

upvoted 4 times

  **elnurgu** 1 year, 8 months ago

Answer C says, "...restrict access to OAI" I obviously didn't understand it. Why do we restrict access to OAI? Actually, I think we need to allow read-only access to OAI.

upvoted 3 times

  **jipark** 10 months, 3 weeks ago

yes, 'restrict' mank confusion.

upvoted 1 times

  **Surferbolt** 1 year, 8 months ago

They probably typed wrongly. If I recall correctly, the button reads somewhere along the lines of 'bucket can restrict access to only OAI', and what it achieves is allowing S3 objects to be accessed only through CloudFront.

upvoted 4 times

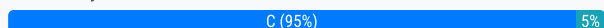
A SysOps administrator is designing a solution for an Amazon RDS for PostgreSQL DB instance. Database credentials must be stored and rotated monthly. The applications that connect to the DB instance send write-intensive traffic with variable client connections that sometimes increase significantly in a short period of time.

Which solution should a SysOps administrator choose to meet these requirements?

- A. Configure AWS Key Management Service (AWS KMS) to automatically rotate the keys for the DB instance. Use RDS Proxy to handle the increases in database connections.
- B. Configure AWS Key Management Service (AWS KMS) to automatically rotate the keys for the DB instance. Use RDS read replicas to handle the increases in database connections.
- C. Configure AWS Secrets Manager to automatically rotate the credentials for the DB instance. Use RDS Proxy to handle the increases in database connections.
- D. Configure AWS Secrets Manager to automatically rotate the credentials for the DB instance. Use RDS read replicas to handle the increases in database connections.

Suggested Answer: C

Community vote distribution



englishborn Highly Voted 1 year, 8 months ago

Selected Answer: C

It cannot be A or B as KMS cannot hold DB credentials, therefore must be Secrets Manager

It cannot be D as the question states "Write intensive" so read only replicas cannot do this.

Has to be C

upvoted 12 times

haxaffee Highly Voted 2 years, 3 months ago

Selected Answer: C

I think C is correct since we are talking about writing to the DB. It is indeed available: Amazon RDS Proxy is available for Amazon Aurora with MySQL compatibility, Amazon Aurora with PostgreSQL compatibility, Amazon RDS for MariaDB, Amazon RDS for MySQL, and Amazon RDS for PostgreSQL.

upvoted 6 times

muznorton Most Recent 6 months, 2 weeks ago

Selected Answer: C

1. Credentials Management

Why AWS Secrets Manager?

AWS Secrets Manager is designed specifically to manage, store, and automatically rotate database credentials securely. It is the most suitable solution for the requirement to rotate credentials monthly for an Amazon RDS instance.

AWS Key Management Service (KMS) is used for encrypting data, not managing database credentials. Thus, KMS does not meet the credentials rotation requirement.

2. Handling Write-Intensive Traffic and Connection Spikes

Why RDS Proxy?

RDS Proxy manages database connections efficiently by pooling and reusing connections, which is critical for applications with variable and spiking client connections.

It helps in reducing the database's memory and CPU overhead, especially during sudden connection surges.

Read replicas are useful for scaling read-only workloads, but they do not assist with write-intensive traffic or connection pooling.

upvoted 2 times

RM999 11 months, 1 week ago

Selected Answer: C

C is correct option

upvoted 1 times

tamng 12 months ago

C is correct

upvoted 1 times

🗨️ 👤 **jipark** 1 year, 4 months ago

Selected Answer: C

why not A :

AWS KMS is used for encryption key management rather than credential rotation.

upvoted 1 times

🗨️ 👤 **Ajit1207** 1 year, 11 months ago

C is right

upvoted 2 times

🗨️ 👤 **michaldavid** 2 years ago

Selected Answer: C

cccccccc

upvoted 2 times

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

Ans: C

upvoted 3 times

🗨️ 👤 **Surferbolt** 2 years, 2 months ago

Selected Answer: C

C.

https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-db.html

upvoted 3 times

🗨️ 👤 **bakjeeone** 2 years, 2 months ago

Selected Answer: C

Secret Manager can rotate secret in database.

upvoted 1 times

🗨️ 👤 **Gianiluca** 2 years, 3 months ago

Selected Answer: A

It's A as Secrets Manager does not auto rotate credentials

upvoted 1 times

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

KMS does not handle database credentials. It handles encryption keys.

This is secrets manager.

upvoted 5 times

A company wants to reduce costs for jobs that can be completed at any time. The jobs currently run by using multiple Amazon EC2 On-Demand Instances and the jobs take slightly less than 2 hours to complete. If a job fails for any reason it must be restarted from the beginning. Which solution will meet these requirements MOST cost-effectively?

- A. Purchase Reserved Instances for the jobs.
- B. Submit a request for a one-time Spot Instance for the jobs.
- C. Submit a request for Spot Instances with a defined duration for the jobs.
- D. Use a mixture of On-Demand Instances and Spot Instances for the jobs.

Suggested Answer: C

Community vote distribution

C (100%)

 **JamesF92**  10 months, 2 weeks ago

Selected Answer: C

It's no longer call "Defined Duration". Now we would use "Persistent Request" not "one-time" request.



<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-requests.html>

Request type: The Spot Instance request type that you choose determines what happens if your Spot Instance is interrupted.

One-time: Amazon EC2 places a one-time request for your Spot Instance. If your Spot Instance is interrupted, the request is not resubmitted.

Persistent request: Amazon EC2 places a persistent request for your Spot Instance. If your Spot Instance is interrupted, the request is resubmitted to replenish the interrupted Spot Instance.

upvoted 6 times

 **kati2k22cz**  1 year, 9 months ago


Selected Answer: C

Its C.

The name of Resource is Spot Block



<https://aws.amazon.com/blogs/aws/new-ec2-spot-blocks-for-defined-duration-workloads/>

upvoted 6 times

 **Surferbolt** 1 year, 8 months ago

Update July 2021 – Spot Instances with a defined duration (also known as Spot blocks) are no longer available to new customers as of July 1, 2021. For customers that have previously used the feature, we will continue to support Spot Instances with a defined duration until December 31, 2022. If your workload is interruption tolerant, we recommend that you use Spot Instances without setting a defined duration. If your workload is not interruption tolerant we recommend that you use On-Demand instances for the required duration of your workload.

upvoted 8 times

 **tuantruong**  3 months, 3 weeks ago

Selected Answer: C


C is right.

upvoted 1 times

 **Liongeek** 1 year, 7 months ago

Until the end of 2022, the ans is C

upvoted 2 times

 **Gomer** 1 year, 2 months ago

Answer will always be C if they don't update the question. The exams are always behind reality (at best).

upvoted 3 times



An environment consists of 100 Amazon EC2 Windows instances. The Amazon CloudWatch agent is deployed and running on all EC2 Instances with a baseline configuration file to capture log files. There is a new requirement to capture the DHCP log files that exist on 50 of the instances. What is the MOST operationally efficient way to meet this new requirement?

- A. Create an additional CloudWatch agent configuration file to capture the DHCP logs. Use the AWS Systems Manager Run Command to restart the CloudWatch agent on each EC2 instance with the append-config option to apply the additional configuration file.
- B. Log in to each EC2 Instance with administrator rights. Create a PowerShell script to push the needed baseline log files and DHCP log files to CloudWatch.
- C. Run the CloudWatch agent configuration file wizard on each EC2 instance. Verify that the baseline log files are included and add the DHCP log files during the wizard creation process.
- D. Run the CloudWatch agent configuration file wizard on each EC2 instance and select the advanced detail level. This will capture the operating system log files.

Suggested Answer: D

Community vote distribution

A (96%) 4%

  **haxaffee** Highly Voted 1 year, 9 months ago

Selected Answer: A

While I find no real solution in the internet for Windows instances, I will go with A because I can't believe logging into EVERY INSTANCE with 100x is a valid solution. Who does that 2022.

upvoted 20 times

  **Hatem08** Most Recent 7 months ago

Selected Answer: A

A is the correct one

upvoted 2 times

  **CVDON** 1 year, 1 month ago

A is more efficient

upvoted 2 times

  **Gomer** 1 year, 1 month ago

Selected Answer: A

Append configuration file (Linux) to running agent so metrics and logs listed in file are collected. (see: "Common scenarios with the CloudWatch agent")

```
/opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-ctl -a append-config -m ec2 -s -c file:/tmp/app.json
```

upvoted 1 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaaa

upvoted 1 times

  **Surferbolt** 1 year, 8 months ago

Selected Answer: A

A. Systems Manager is the way to go.


upvoted 3 times

  **sassy69** 1 year, 8 months ago

I go with A.

D is valid, but not the MOST operationally efficient.

upvoted 2 times

  **kati2k22cz** 1 year, 9 months ago

Selected Answer: D

D.

More info of detail level here

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/create-cloudwatch-agent-configuration-file-wizard.html>

upvoted 1 times

  **Liongeek** 1 year, 7 months ago

The wizard only can be used if you log to the instance. But there are 100 instances. That solution doesn't apply.

upvoted 4 times

A company has 10 Amazon EC2 instances in its production account. A SysOps administrator must ensure that email notifications are sent to administrators each time there is an EC2 instance state change.

Which solution will meet this requirements?

- A. Configure an Amazon Route 53 simple routing policy that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic when an EC2 instance state changes. This SNS topic then sends notifications to its email subscribers.
- B. Configure an Amazon Route 53 simple routing policy that publishes a message to an Amazon Simple Queue Service (Amazon SQS) queue when an EC2 instance state changes. This SQS queue then sends notifications to its email subscribers.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic when an EC2 instance state changes. This SNS topic then sends notifications to its email subscribers.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that publishes a message to an Amazon Simple Queue Service (Amazon SQS) queue when an EC2 instance state changes. This SQS queue then sends notifications to its email subscribers.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **tamng** 12 months ago

C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic when an EC2 instance state changes. This SNS topic then sends notifications to its email subscribers.

upvoted 2 times

🗳️ 👤 **ronnykapo** 1 year, 6 months ago

CCCCCCCC

upvoted 4 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: C

CCCCCCCC

upvoted 2 times

🗳️ 👤 **Surferbolt** 2 years, 2 months ago

Selected Answer: C

C. EventBridge and SNS.

upvoted 3 times

🗳️ 👤 **XAvenge** 2 years, 3 months ago

Selected Answer: C

C. All Route 53 and SQS options are automatically out. C is left

upvoted 4 times

A company has an application that runs on a fleet of Amazon EC2 instances behind an Elastic Load Balancer. The instances run in an Auto Scaling group. The application's performance remains consistent throughout most of each day. However, an increase in user traffic slows the performance during the same 4-hour period of time each day.


What is the MOST operationally efficient solution that will resolve this issue?

- A. Configure a second Elastic Load Balancer in front of the Auto Scaling group with a weighted routing policy.
- B. Configure the fleet of EC2 instances to run on larger instance types to support the increase in user traffic.
- C. Create a scheduled scaling action to scale out the number of EC2 instances shortly before the increase in user traffic occurs.
- D. Manually add a few more EC2 instances to the Auto Scaling group to support the increase in user traffic.

Suggested Answer: A

Community vote distribution

C (100%)

 **princajen** Highly Voted 2 years, 4 months ago

Selected Answer: C

I'm voting for C!

They see the same slow performance at the same time each day. Scheduled scaling makes sense.

Scheduled scaling helps you to set up your own scaling schedule according to predictable load changes. For example, let's say that every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can configure a schedule for Amazon EC2 Auto Scaling to increase capacity on Wednesday and decrease capacity on Friday.

To use scheduled scaling, you create scheduled actions. Scheduled actions are performed automatically as a function of date and time. When you create a scheduled action, you specify when the scaling activity should occur and the new desired, minimum, and maximum sizes for the scaling action. You can create scheduled actions that scale one time only or that scale on a recurring schedule.

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

upvoted 13 times

 **haxaffee** Highly Voted 2 years, 3 months ago

Selected Answer: C

Vote C. No way its A.

upvoted 5 times

 **stoy123** Most Recent 10 months, 1 week ago

Selected Answer: C

C it is

upvoted 2 times

 **tamng** 12 months ago

vOTE C

upvoted 1 times

 **BrascoChe** 1 year, 3 months ago

I very much think it was C because schedule action here makes a lot of sense.

upvoted 2 times

 **fazlur21** 1 year, 6 months ago

Selected Answer: C

C the Answer

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

upvoted 3 times

 **michaldavid** 2 years ago

Selected Answer: C

ccccccc

upvoted 4 times

  **Liongeek** 2 years, 1 month ago

Ans: C



upvoted 1 times

  **Surferbolt** 2 years, 2 months ago

Selected Answer: C

C. Since peak traffic is predictable, they can schedule a scale out.

upvoted 3 times

  **Rick365** 2 years, 3 months ago

Selected Answer: C

C. Create a scheduled scaling action to scale out the number of EC2 instances shortly before the increase in user traffic occurs.

upvoted 3 times

A company hosts an application on an Amazon EC2 instance in a single AWS Region. The application requires support for non-HTTP TCP traffic and HTTP traffic.

The company wants to deliver content with low latency by leveraging the AWS network. The company also wants to implement an Auto Scaling group with an Elastic Load Balancer.

How should a SysOps administrator meet these requirements?


- A. Create an Auto Scaling group with an Application Load Balancer (ALB). Add an Amazon CloudFront distribution with the ALB as the origin.
- B. Create an Auto Scaling group with an Application Load Balancer (ALB). Add an accelerator with AWS Global Accelerator with the ALB as an endpoint.
- C. Create an Auto Scaling group with a Network Load Balancer (NLB). Add an Amazon CloudFront distribution with the NLB as the origin.
- D. Create an Auto Scaling group with a Network Load Balancer (NLB). Add an accelerator with AWS Global Accelerator with the NLB as an endpoint.

Suggested Answer: B

Community vote distribution

D (88%)

6%

 **Gorille69** Highly Voted 2 years, 3 months ago

It is not A or B because ALB works with layer 7 of OSI model (HTTP, HTTPS) , NLB works with layer 4 (TCP, UDP, TLS) . it remains C or D . it is possible to have the eLB as the origin Cloudfront but CloudFront is designed to handle the HTTP protocol, while Global Accelerator is best used for HTTP and non-HTTP protocols such as TCP and UDP. So, my answer is D too.

upvoted 23 times

 **haxaffee** Highly Voted 2 years, 3 months ago

Selected Answer: D

In my opinion every answer so far is wrong. The question states it needs HTTP and NON-HTTP TCP traffic. Therefore it cannot be an ALB and must be a NLB. AWS Global Accelerator: "When the internet is congested, AWS Global Accelerator optimizes the path to your application to keep packet loss, jitter, and latency consistently low."

D is correct in my world.

upvoted 10 times

 **Grodgar** Most Recent 6 months, 1 week ago

Selected Answer: D

The problem is the question is:

- 1) non-http traffic --> Global accelerator
- 2) in a single Region --> Cloud Front.

All answers are not correct, but exam wants us to evaluate traffic protocols over locations

upvoted 1 times

 **stoy123** 10 months, 1 week ago

Selected Answer: D

D it is

upvoted 2 times

 **tamng** 12 months ago

D. Create an Auto Scaling group with a Network Load Balancer (NLB). Add an accelerator with AWS Global Accelerator with the NLB as an endpoint.

- TCP (Layer 4) => NLB

- CloudFront + ALB : ok, but CloudFront + NLB: Never => AWS Global Accelerator with the NLB as an endpoint : good mix

=> D

upvoted 2 times

 **mussaha** 1 year, 2 months ago

Selected Answer: D

NLB Deals with TCP

Cloudfront cannot deal with non-HTTP traffic

upvoted 3 times

🗳️ 👤 **callspace** 1 year, 2 months ago

Selected Answer: D

HTTP/TCP and low latency by leveraging the AWS network. Option D has endpoint as a clue!

upvoted 2 times

🗳️ 👤 **sxti** 1 year, 6 months ago

The Application Load Balancer (ALB) is a good choice for HTTP and HTTPS traffic, but it doesn't support non-HTTP TCP traffic. So, we can't choose options A and B.

Network Load Balancer (NLB) supports load balancing of both TCP traffic and HTTP traffic, so we can use it in this case.

The Amazon CloudFront distribution is primarily for caching static and dynamic content closer to viewers to reduce latency, and it doesn't directly support non-HTTP protocols or load balancing for EC2 instances. So, we can't choose option C.

upvoted 3 times

🗳️ 👤 **CVDON** 1 year, 7 months ago

D because you need support for http and non-http protocols

upvoted 1 times

🗳️ 👤 **Gomer** 1 year, 8 months ago

Selected Answer: D

Without question, answer can only be D. The table comparison in the following link makes that perfectly clear: <https://jayendrapatil.com/aws-cloudfront-vs-global-accelerator/>

upvoted 2 times

🗳️ 👤 **AMYadav** 1 year, 9 months ago

Selected Answer: B

As per AAWS Global Accelerator FAQs, 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 1 times

🗳️ 👤 **braveheart22** 1 year, 10 months ago

DDD is the correct option from my point of view. The key phrase here is "The application requires support for non-HTTP TCP traffic and HTTP traffic" and as we know, HTTP is layer 7 of the OSI model (Application layer). TCP and UDP are layer 4 and therefore it must be Network Load Balancer (NLB). Global Accelerator is a good fit for non-HTTP use cases, which is the reason that I'm inclined towards D.

upvoted 2 times

🗳️ 👤 **Akaza** 2 years ago

D for sure

upvoted 3 times

🗳️ 👤 **Surferbolt** 2 years, 2 months ago

Selected Answer: D

D is the answer.

upvoted 4 times

🗳️ 👤 **Balliache520505** 2 years, 3 months ago

Answer D. AWS Global Accelerator and Amazon CloudFront are separate services that use the AWS global network and its edge locations around the world. CloudFront improves performance for both cacheable content (such as images and videos) and dynamic content (such as API acceleration and dynamic site delivery). 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 4 times

🗳️ 👤 **kati2k22cz** 2 years, 3 months ago

Selected Answer: D

Yes, found 2 references. I agree with D.

<https://medium.com/awesome-cloud/aws-difference-between-application-load-balancer-and-network-load-balancer-cb8b6cd296a4>

https://aws.amazon.com/global-accelerator/faqs/?nc1=h_ls

upvoted 4 times

  **Lolaadmin** 2 years, 3 months ago

Ans:D

CloudFront is designed to handle HTTP protocol meanwhile Global Accelerator is best used for both HTTP and non-HTTP protocols such as TCP and UDP.

<https://tutorialsdojo.com/aws-global-accelerator-vs-amazon-cloudfront/>

upvoted 1 times

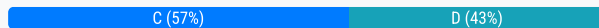
A SysOps administrator has an AWS CloudFormation template that is used to deploy an encrypted Amazon Machine Image (AMI). The CloudFormation template will be used in a second account so the SysOps administrator copies the encrypted AMI to the second account. When launching the new CloudFormation stack in the second account, it fails.

Which action should the SysOps administrator take to correct the issue?

- A. Change the AMI permissions to mark the AMI as public.
- B. Deregister the AMI in the source account.
- C. Re-encrypt the destination AMI with an AWS Key Management Service (AWS KMS) key from the destination account.
- D. Update the CloudFormation template with the ID of the AMI in the destination account.

Suggested Answer: C

Community vote distribution



princajen Highly Voted 2 years, 9 months ago

Selected Answer: C

C!

While launching the instance from a shared encrypted AMI, you can specify a KMS key of your choice. You may also choose cmkSource to encrypt volumes in your account. However, we recommend that you re-encrypt the volumes using a KMS key in the target account. This protects you if the source KMS key is compromised, or if the source account revokes permissions, which could cause you to lose access to any encrypted volumes you created using cmkSource.

<https://aws.amazon.com/blogs/security/how-to-share-encrypted-amis-across-accounts-to-launch-encrypted-ec2-instances/>
upvoted 10 times

TareDHakim 1 year, 6 months ago

thanks for the URL - however, this URL indicates the answer is actually D and NOT C.
upvoted 2 times

braveheart22 Highly Voted 2 years, 4 months ago

I'm inclined toward D in every sense.
upvoted 7 times

tamng 1 year, 6 months ago

C is correct
upvoted 1 times

heshamahammad Most Recent 2 months ago

Selected Answer: C

Definetly C, "While launching the instance from a shared encrypted AMI, you can specify a KMS key of your choice. You may also choose cmkSource to encrypt volumes in your account. However, we recommend that you re-encrypt the volumes using a KMS key in the target account. This protects you if the source KMS key is compromised, or if the source account revokes permissions, which could cause you to lose access to any encrypted volumes you created using cmkSource."

<https://aws.amazon.com/blogs/security/how-to-share-encrypted-amis-across-accounts-to-launch-encrypted-ec2-instances/>
upvoted 1 times

arsovai 2 months ago

Selected Answer: C

Even if you update the AMI ID, if the encryption issue isn't resolved, the stack will still fail
upvoted 1 times

OlehKom 6 months, 2 weeks ago

Selected Answer: C

D only can be a valid solution if it explicitly says that the 2d has access to the shared AMI and KSM key. Only after that you could update the template with the existing AMI id

BUT it's not explicitly mentioned that the second account has been granted access to the encrypted AMI or the KMS key. Without explicitly granting access, the second account cannot use the AMI. This is why re-encryption in the second account can be a necessary step to ensure smooth deployment, so C

upvoted 1 times

🗳️ 👤 **numark** 7 months, 1 week ago

Answer is D: When an encrypted Amazon Machine Image (AMI) is copied to another AWS account, it gets a new AMI ID in the destination account. The CloudFormation template in the second account must reference the new AMI ID to successfully launch the stack.

Not C: While this step is often part of copying an encrypted AMI between accounts, it does not fix the issue in this scenario. The new AMI in the destination account is already encrypted with a KMS key from the second account. The failure occurs because the CloudFormation template is still referencing the old AMI ID.

upvoted 1 times

🗳️ 👤 **tsangckl** 1 year, 1 month ago

Selected Answer: D

By copilot

The SysOps administrator should take action D. Update the CloudFormation template with the ID of the AMI in the destination account.

When an AMI is copied to another account, it gets a new AMI ID. The CloudFormation template in the second account is likely still referencing the AMI ID from the original account, which is causing the stack deployment to fail. By updating the template with the new AMI ID, the stack deployment should proceed without issues. Other options (A, B, C) are not relevant to the issue described.

upvoted 4 times

🗳️ 👤 **tsangckl** 1 year, 1 month ago

When you copy an AMI, the copied AMI is encrypted using the same AWS KMS key as the original AMI, by default. If the original AMI is encrypted with a default EBS key, the copied AMI will also be encrypted with a default EBS key, and this key will be unique to the account to which the AMI is copied. Therefore, there's no need to re-encrypt the AMI in the destination account. The issue here is that the CloudFormation template is still referencing the old AMI ID from the source account, which is not recognized in the destination account. Hence, updating the CloudFormation template with the new AMI ID (option D) is the correct action to resolve the issue.

upvoted 3 times

🗳️ 👤 **AgboolaKun** 1 year, 1 month ago

Selected Answer: C

Sincerely, any of the C or D could cause the template to fail. However, I will go for C since the emphasis is on encryption and best practice in this situation expects that you re-encrypt the volumes using a KMS key in the destination account.

upvoted 1 times

🗳️ 👤 **a6a3d55** 1 year, 2 months ago

Selected Answer: D

Even if you reencrypt it will still not work until the AMI ID is changed in the template

upvoted 3 times

🗳️ 👤 **seetpt** 1 year, 2 months ago

Selected Answer: C

I vote for C

upvoted 1 times

🗳️ 👤 **stoy123** 1 year, 4 months ago

Selected Answer: D

D for sure

upvoted 4 times

🗳️ 👤 **henro4niger** 1 year, 4 months ago

I will go with D. Why do I need to re-encrypt the ami in the second account when I can just update the template with the ami ID of the target account?

D is definitely the answer, moreover, C will introduce serious complexity

upvoted 1 times

🗳️ 👤 **vivanchyk** 1 year, 3 months ago

While it is necessary to update the CloudFormation template with the new AMI ID after copying it to the destination account, this action alone won't solve the encryption key access issue. The failure is likely due to the lack of access to the KMS key, not merely the AMI ID reference

When an AMI is encrypted, it is done so using a specific AWS Key Management Service (KMS) key. If you copy an encrypted AMI to another AWS

account, the destination account needs appropriate permissions to use the KMS key that encrypted the AMI, or the AMI needs to be re-encrypted with a KMS key that belongs to the destination account.

upvoted 1 times

🗨️ 👤 **xdkonorek2** 1 year, 5 months ago

Selected Answer: D

Definitely D

Copied AMI already is encrypted by KMS key that is stored in target aws account.

upvoted 2 times

🗨️ 👤 **vivanchyk** 1 year, 3 months ago

where is it said that "Copied AMI already is encrypted by KMS key that is stored in target aws account." ???

upvoted 1 times

🗨️ 👤 **Aamee** 9 months, 2 weeks ago

See this comment by a user above. This is what he said to justify the ans. D selection is correct here:

"When you copy an AMI, the copied AMI is encrypted using the same AWS KMS key as the original AMI, by default. If the original AMI is encrypted with a default EBS key, the copied AMI will also be encrypted with a default EBS key, and this key will be unique to the account to which the AMI is copied. "

upvoted 1 times

🗨️ 👤 **tamng** 1 year, 6 months ago

C not D

upvoted 1 times

🗨️ 👤 **Hatem08** 1 year, 6 months ago

Selected Answer: C

ccccccc

upvoted 2 times

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

Answer is C

upvoted 2 times

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

Selected Answer: D

"Copying a source AMI results in an identical but distinct target AMI with its own unique identifier. You can change or deregister the source AMI with no effect on the target AMI. The reverse is also true."

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/CopyingAMIs.html>

upvoted 5 times

A company's SysOps administrator deploys four new Amazon EC2 instances by using the standard Amazon Linux 2 Amazon Machine Image (AMI). The company needs to be able to use AWS Systems Manager to manage the instances. The SysOps administrator notices that the instances do not appear in the Systems Manager console.

What must the SysOps administrator do to resolve this issue?

- A. Connect to each instance by using SSH. Install Systems Manager Agent on each instance. Configure Systems Manager Agent to start automatically when the instances start up.
- B. Use AWS Certificate Manager (ACM) to create a TLS certificate. Import the certificate into each instance. Configure Systems Manager Agent to use the TLS certificate for secure communications.
- C. Connect to each instance by using SSH. Create an ssm-user account. Add the ssm-user account to the /etc/sudoers.d directory.
- D. Attach an IAM instance profile to the instances. Ensure that the instance profile contains the AmazonSSMManagedInstanceCore policy.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Liongeek** Highly Voted 2 years, 7 months ago

Ans: D

If instances were deployed with Amazon Linux 2, they already have the SSM Agent installed.

upvoted 8 times

🗳️ 👤 **Xelnak** Highly Voted 2 years, 7 months ago

D. for sure

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

upvoted 5 times

🗳️ 👤 **numark** Most Recent 6 months, 3 weeks ago

Selected Answer: D

To use AWS Systems Manager to manage EC2 instances, each instance must have an IAM role with the necessary permissions attached to it. This role is known as an instance profile. The AmazonSSMManagedInstanceCore policy grants the necessary permissions for Systems Manager to interact with the EC2 instances.

upvoted 1 times

🗳️ 👤 **tamng** 1 year, 6 months ago

D is correct

upvoted 1 times

🗳️ 👤 **shubham2705** 2 years, 2 months ago

Its A.

To enable AWS Systems Manager to manage EC2 instances, the Systems Manager Agent must be installed and running on each instance. The Systems Manager Agent is responsible for sending data about the instance to AWS Systems Manager and executing commands and scripts on the instance as part of Systems Manager features such as Run Command, Patch Manager, and Session Manager.

Option D is incorrect as it refers to attaching an IAM instance profile to the instances with the AmazonSSMManagedInstanceCore policy, which is not necessary for resolving the issue of instances not appearing in the Systems Manager console.

upvoted 2 times

🗳️ 👤 **apreda** 2 years ago

It's D. AWS System Manager already installed on Amazon Linux OS.

upvoted 4 times

🗳️ 👤 **fazlur21** 2 years ago

"Amazon ECS Optimized Linux 2 Amazon Machine Images(AMI) now come with the AWS System Manager (SSM) Agent pre-installed"

ref : [https://aws.amazon.com/about-aws/whats-new/2020/02/amazon-ecs-optimized-linux-2-amis-come-pre-installed-aws-systems-manager-agent/#:~:text=Amazon%20ECS%20Optimized%20Linux%20,SSM\)%20Agent%20pre%2Dinstalled.](https://aws.amazon.com/about-aws/whats-new/2020/02/amazon-ecs-optimized-linux-2-amis-come-pre-installed-aws-systems-manager-agent/#:~:text=Amazon%20ECS%20Optimized%20Linux%20,SSM)%20Agent%20pre%2Dinstalled.)

upvoted 3 times

  **braveheart22** 2 years, 4 months ago

DDDDDDD it is.

upvoted 2 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddd

upvoted 2 times

A SysOps administrator is maintaining a web application using an Amazon CloudFront web distribution, an Application Load Balancer (ALB), Amazon RDS, and Amazon EC2 in a VPC. All services have logging enabled. The administrator needs to investigate HTTP Layer 7 status codes from the web application.

Which log sources contain the status codes? (Choose two.)

- A. VPC Flow Logs
- B. AWS CloudTrail logs
- C. ALB access logs
- D. CloudFront access logs
- E. RDS logs

Suggested Answer: CD



Community vote distribution

CD (100%)

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

Selected Answer: CD

level 7 -> ALB + CloudFront
upvoted 7 times

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

Ans: C,D
upvoted 5 times

  **jjpark** Most Recent 10 months, 3 weeks ago

Selected Answer: CD


why not B :
AWS CloudTrail logs record API calls and events for AWS services, but they are not focused on capturing HTTP Layer 7 status codes from web application traffic.
upvoted 5 times

  **McEgowan** 11 months, 3 weeks ago

I have never heard of CloudFront access logs in AWS so that must be wrong and i wonder everyone chose it?
upvoted 2 times

  **KTsankov** 11 months, 1 week ago

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/AccessLogs.html>
upvoted 3 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: CD

C and D
upvoted 1 times

  **goatbernard** 1 year, 6 months ago

level 7 -> ALB + CloudFront
upvoted 3 times

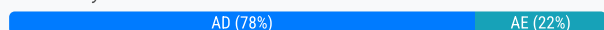
A company wants to be alerted through email when IAM CreateUser API calls are made within its AWS account.

Which combination of actions should a SysOps administrator take to meet this requirement? (Choose two.)

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with AWS CloudTrail as the event source and IAM CreateUser as the specific API call for the event pattern.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with Amazon CloudSearch as the event source and IAM CreateUser as the specific API call for the event pattern.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with AWS IAM Access Analyzer as the event source and IAM CreateUser as the specific API call for the event pattern.
- D. Use an Amazon Simple Notification Service (Amazon SNS) topic as an event target with an email subscription.
- E. Use an Amazon Simple Email Service (Amazon SES) notification as an event target with an email subscription.

Suggested Answer: AD

Community vote distribution



michaldavid Highly Voted 2 years ago

Selected Answer: AD

A and D

upvoted 6 times

Liongeek Highly Voted 2 years, 1 month ago

Ans: A,D

upvoted 5 times

Rado_Piatek Most Recent 1 month, 2 weeks ago

Selected Answer: AD

A,D

not E: SES is used to send emails to customers (transactions etc.), not for alerting

upvoted 1 times

tamng 12 months ago

A and D

upvoted 2 times

jipark 1 year, 4 months ago

Selected Answer: AE

why not D :

Amazon SES (Simple Email Service) can send emails, but it is not typically used as a direct event target for CloudWatch Events.

upvoted 3 times

jipark 1 year, 4 months ago

my mistype (please remove first comment :

why not E :

Amazon SES (Simple Email Service) can send emails, but it is not typically used as a direct event target for CloudWatch Events

upvoted 2 times

fazlur21 1 year, 6 months ago

Selected Answer: AD

A and D the Answer

upvoted 3 times

A database is running on an Amazon RDS Multi-AZ DB instance. A recent security audit found the database to be out of compliance because it was not encrypted.

Which approach will resolve the encryption requirement?

- A. Log in to the RDS console and select the encryption box to encrypt the database.
- B. Create a new encrypted Amazon EBS volume and attach it to the instance.
- C. Encrypt the standby replica in the secondary Availability Zone and promote it to the primary instance.
- D. Take a snapshot of the RDS instance, copy and encrypt the snapshot, and then restore to the new RDS instance.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Xelnak** Highly Voted 👍 7 months, 2 weeks ago

D. You can only encrypt an Amazon RDS DB instance when you create it, not after the DB instance is created.

However, because you can encrypt a copy of an unencrypted snapshot, you can effectively add encryption to an unencrypted DB instance. That is, you can create a snapshot of your DB instance, and then create an encrypted copy of that snapshot. You can then restore a DB instance from the encrypted snapshot, and thus you have an encrypted copy of your original DB instance.

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

upvoted 9 times

🗳️ 👤 **michaldavid** Most Recent 🕒 6 months, 3 weeks ago

Selected Answer: D

dddddddddd

upvoted 2 times

🗳️ 👤 **Liongeek** 7 months, 2 weeks ago

Selected Answer: D

Ans: D

Ref: <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html#Overview.Encryption.Limitations>

upvoted 3 times

A company using AWS Organizations requires that no Amazon S3 buckets in its production accounts should ever be deleted.



What is the SIMPLEST approach the SysOps administrator can take to ensure S3 buckets in those accounts can never be deleted?

- A. Set up MFA Delete on all the S3 buckets to prevent the buckets from being deleted.
- B. Use service control policies to deny the s3:DeleteBucket action on all buckets in production accounts.
- C. Create an IAM group that has an IAM policy to deny the s3:DeleteBucket action on all buckets in production accounts.
- D. Use AWS Shield to deny the s3:DeleteBucket action on the AWS account instead of all S3 buckets.

Suggested Answer: B

Community vote distribution



B (100%)

  **michaldavid** Highly Voted 2 years, 6 months ago

Selected Answer: B

bbbbbbb

upvoted 6 times

  **jipark** 1 year, 10 months ago

SCP deny user's action

upvoted 3 times

  **noircesar25** Most Recent 1 year, 1 month ago



you cant delete a bucket if its not empty. so enable MFA would solve the problem because only the root user can delete objects

upvoted 1 times

  **Aamee** 9 months, 2 weeks ago

The question specifically asks for the production account S3 buckets. The MFA option suggests that it applies to all the buckets regardless of any accounts it currently reside in. Plus, it's not operationally efficient to apply the MFA option across several S3 buckets manually compared to just configure through SCP policy across the Org level and for all Production accounts automatically.

upvoted 1 times

  **jipark** 1 year, 10 months ago

Selected Answer: B

why not C :

assigning IAM group to each user manually is not efficient.

upvoted 3 times

  **Liongeek** 2 years, 7 months ago

Ans: B

upvoted 4 times

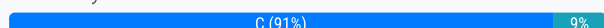
A company has an application that is running on Amazon EC2 instances in a VPC. The application needs access to download software updates from the internet. The VPC has public subnets and private subnets. The company's security policy requires all EC2 instances to be deployed in private subnets.

What should a SysOps administrator do to meet these requirements?

- A. Add an internet gateway to the VPC. In the route table for the private subnets, add a route to the internet gateway.
- B. Add a NAT gateway to a private subnet. In the route table for the private subnets, add a route to the NAT gateway.
- C. Add a NAT gateway to public subnet. In the route table for the private subnets, add a route to the NAT gateway.
- D. Add two internet gateways to the VPC. In the route tables for the private subnets and public subnets, add a route to each internet gateway.

Suggested Answer: C

Community vote distribution



sxti Highly Voted 2 years ago

Selected Answer: C

The correct answer is C. Add a NAT gateway to a public subnet. In the route table for the private subnets, add a route to the NAT gateway.

The application needs to be able to download updates from the internet, but it's running on EC2 instances in a private subnet. Private subnets do not have direct access to the internet. A NAT gateway allows instances in a private subnet to connect to the internet or other AWS services but prevent the internet from initiating a connection with those instances.

upvoted 6 times

OlehKom Most Recent 6 months, 2 weeks ago

Selected Answer: C

Internet Gateway: Public access (e.g., web server).

NAT Gateway: Private access to the internet (e.g., backend or database servers)

upvoted 1 times

10cc6ba 11 months, 3 weeks ago

Selected Answer: C

C is correct

upvoted 1 times

Rabbit117 1 year, 3 months ago

Selected Answer: C

C is correct. The NAT Gateway is deployed in the Public subnet and the route table for the private subnet points all internet bound traffic to the NAT GW.

upvoted 2 times

stoy123 1 year, 4 months ago

Selected Answer: C

C, NAT Gateway must be created in a public subnet

upvoted 3 times

JamesF92 1 year, 10 months ago

Selected Answer: C

<https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-scenarios.html>

This one is definitely C.

upvoted 4 times

fazlur21 2 years ago

B

why? because a public subnet don't need nat gateway, only private subnets need a nat gateway to connect to the internet

upvoted 1 times

🗨️ **wookchan** 1 year, 11 months ago

No. The NAT gateway must be deployed in a public subnet, because it needs to be connected to the Internet Gateway.

upvoted 3 times

🗨️ **vinothc** 2 years, 2 months ago

To enable the EC2 instances in private subnets to download software updates from the internet, a SysOps administrator should add a NAT gateway to a private subnet, and in the route table for the private subnets, add a route to the NAT gateway. Therefore, option B is the correct answer.

Option A is incorrect because adding an internet gateway to the VPC and a route to the internet gateway in the private subnet's route table would not work since the private subnet does not have a public IP address.

Option C is incorrect because adding a NAT gateway to the public subnet and a route to the NAT gateway in the private subnet's route table would not work because the private subnet requires outbound traffic to traverse the NAT gateway, which would be difficult to implement in a security perspective.

Option D is incorrect because having two internet gateways is not practical and would not resolve the issue of allowing private instances to download software updates from the internet.

upvoted 2 times

🗨️ **henryford** 2 years, 1 month ago

This is why I'm not afraid that ChatGPT will replace us soon. Option B is clearly wrong as your NAT gateway must have a public IP. The reasoning behind why Option C is also clearly incorrect as a route in the routing table would obviously cause the traffic to be routed through the NAT gateway.

upvoted 5 times

🗨️ **Gomer** 2 years, 2 months ago

Selected Answer: C

Can't be D. Can have only one IGW per VPC. Need IGW and route to a NAT Gateway from private subnet.

upvoted 2 times

🗨️ **skywalker** 2 years, 5 months ago

Selected Answer: C

CCCCCCCC

upvoted 4 times

🗨️ **jessbase2022** 2 years, 6 months ago

Selected Answer: C

C

<https://aws.amazon.com/premiumsupport/knowledge-center/nat-gateway-vpc-private-subnet/>

upvoted 4 times

🗨️ **MrMLB** 2 years, 6 months ago

Selected Answer: B

To meet the requirements of the company's security policy, the SysOps administrator should choose option B: Add a NAT gateway to a private subnet. In the route table for the private subnets, add a route to the NAT gateway.

In this scenario, the EC2 instances in the private subnets need access to the internet to download software updates, but they cannot be directly connected to the internet. A NAT gateway allows the EC2 instances to connect to the internet indirectly by routing their traffic through the NAT gateway, which is located in a public subnet. The NAT gateway has a direct connection to the internet, so it can access the internet on behalf of the EC2 instances.

To set this up, the SysOps administrator should create a NAT gateway in a private subnet and then add a route to the NAT gateway in the route table for the private subnets. This will allow the EC2 instances in the private subnets to access the internet through the NAT gateway.

upvoted 3 times

🗨️ **henryford** 2 years, 1 month ago

You can't use a NAT gateway in a private subnet as it needs a public IP address to work.



upvoted 4 times

🗨️ **michaldavid** 2 years, 6 months ago

Selected Answer: C

cccccccc

upvoted 3 times

  **tts1234** 2 years, 6 months ago

Selected Answer: C

A Nat Gateway enables instances in private subnets to connect to the internet. The Nat gateway must be deployed in the public subnet with an Elastic IP. Once the resource is created, a route table associated with the private subnet needs to point internet-bound traffic to the NAT gateway.

<https://towardsdatascience.com/connecting-to-an-ec2-instance-in-a-private-subnet-on-aws-38a3b86f58fb>

upvoted 3 times

A development team recently deployed a new version of a web application to production. After the release, penetration testing revealed a cross-site scripting vulnerability that could expose user data.

Which AWS service will mitigate this issue?

- A. AWS Shield Standard
- B. AWS WAF
- C. Elastic Load Balancing
- D. Amazon Cognito

Suggested Answer: B

Community vote distribution

B (100%)

  **bogossdu35** Highly Voted 2 years, 2 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-xss-match.html>



upvoted 8 times

  **apmmahesh** Most Recent 6 months ago

Selected Answer: B

AWS WAF is a security service that helps protect web applications from common threats such as SQL injection and cross-site scripting.



upvoted 1 times

  **10cc6ba** 11 months, 3 weeks ago

Selected Answer: B

B of course

upvoted 1 times



  **Rabbit117** 1 year, 3 months ago

Selected Answer: B

B.

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-xss-match.html>

upvoted 2 times

  **tamng** 1 year, 6 months ago

B is correct

upvoted 2 times

  **strovertz** 1 year, 7 months ago

Selected Answer: B

B

read <https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-xss-match.html>

upvoted 3 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbb

upvoted 2 times

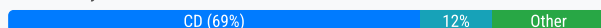
A SysOps administrator must configure a resilient tier of Amazon EC2 instances for a high performance computing (HPC) application. The HPC application requires minimum latency between nodes.

Which actions should the SysOps administrator take to meet these requirements? (Choose two.)

- A. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the file system to the EC2 instances by using user data.
- B. Create a Multi-AZ Network Load Balancer in front of the EC2 instances.
- C. Place the EC2 instances in an Auto Scaling group within a single subnet.
- D. Launch the EC2 instances into a cluster placement group.
- E. Launch the EC2 instances into a partition placement group.

Suggested Answer: CD

Community vote distribution



Gomer Highly Voted 1 year, 8 months ago

Selected Answer: CD

CD

D. Amazon ALWAYS touts the "cluster placement group" as being the solution HPC workloads (low latency between nodes). However this actually precludes resilience through ELB and multi-AZ spread of instances or multiple partitions (insufficient throughput).

C. The only option I see to get both resilience and HPC is to rely on auto-scaling group to re-scale for any failing nodes. Provides some protection unless the whole rack fails. No, it's not multi-AZ or multi rack, but its better than nothing given Amazons HPC recommendations.

upvoted 9 times

Vivec Highly Voted 1 year, 9 months ago

Selected Answer: CD

To meet the requirement of minimum latency between nodes for a high performance computing (HPC) application on a resilient tier of Amazon EC2 instances, the SysOps administrator should take the following actions:

Launch the EC2 instances into a cluster placement group to ensure that the instances are placed in a low-latency, single-tenant infrastructure.

Place the EC2 instances in an Auto Scaling group within a single subnet. This will ensure that the instances are in the same Availability Zone and will not have to traverse a network boundary to communicate with each other.

upvoted 7 times

Student013657 Most Recent 6 months, 3 weeks ago

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

upvoted 1 times

Rabbit117 9 months, 3 weeks ago

Selected Answer: CD

C and D.

An Auto Scaling group will provide the resilience and the Cluster placement group will provide the low latency.

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

upvoted 3 times

Yowie351 10 months ago

Selected Answer: BE

B and E

upvoted 1 times

tamng 12 months ago

C. Place the EC2 instances in an Auto Scaling group within a single subnet.

D. Launch the EC2 instances into a cluster placement group.

upvoted 1 times

🗳️ 👤 **wh1t4k3r** 1 year, 3 months ago

I was going for B and D too, but here's the thing with B:

Cluster placement groups are deployed within AN availability zone, meaning the multi-az LB does not apply:

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

upvoted 1 times

🗳️ 👤 **tamng** 12 months ago

C and D, not B. You wrong

upvoted 1 times

🗳️ 👤 **kondratyevmn** 1 year, 7 months ago

Selected Answer: BE

E - you need resilient tier, which partition placement groups provide, not C because it doesn't (although it provides minimum latency between nodes).

B - you need LB with low latency to distribute traffic between your instances, placed in partition placement groups.

upvoted 2 times

🗳️ 👤 **wooyourdaddy** 1 year, 8 months ago

Selected Answer: CD

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

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

upvoted 1 times

🗳️ 👤 **nearavenac** 1 year, 10 months ago

Selected Answer: AD

ADADAD

upvoted 1 times

🗳️ 👤 **csG13** 1 year, 10 months ago

Selected Answer: AD

I vote for A & D. Question is about low latency and HPC, therefore D is the right placement option (so, no E). Furthermore, cluster placement will use a single availability zone, hence multi-AZ NLB is useless (so, no B). Finally, regarding ASG in a single subnet isn't resilient. Below I paste directly from AWS docs:

"If you try to add more instances to the placement group later, or if you try to launch more than one instance type in the placement group, you increase your chances of getting an insufficient capacity error."

So in case of a failed health check, ASG might try to spin up a new instance but receive an `InsufficientInstanceCapacity` (which is an AZ-specific error). Therefore, is not recommended when resiliency is needed. The only thing left is A.

PS.

Not a very good question, I believe a bit more context is required.

upvoted 4 times

🗳️ 👤 **Untamables** 1 year, 10 months ago

Selected Answer: BD

B for resilience

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/disaster-recovery-resiliency.html>

D for minimum latency

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

upvoted 1 times

🗳️ 👤 **noahsark** 1 year, 11 months ago

Selected Answer: CE

Place the EC2 instances in an Auto Scaling group within a single subnet.

Launch the EC2 instances into a partition placement group.

Not B

While Multi-AZ deployment offers high availability, some workloads are more sensitive to internode latency and could not be deployed across multiple zones. With partition placement groups, you can now deploy these workloads within a single zone and reduce the likelihood of correlated failures,

improving your application performance and availability.

Not D

Customers wanted a way to reduce correlated failures for large distributed and replicated workloads that required hundreds of EC2 instances.

To isolate the impact of hardware faults, EC2 subdivides each partition placement group, into logical segments called partitions. EC2 ensures that no two partitions within a placement group share the same racks.

<https://aws.amazon.com/blogs/compute/using-partition-placement-groups-for-large-distributed-and-replicated-workloads-in-amazon-ec2/>
upvoted 2 times

🗨️ 👤 **Spike2020** 1 year, 11 months ago

BE

Resilient is not 1 subnet. Hence a network load balancer.

upvoted 1 times

🗨️ 👤 **Polietylen** 1 year, 11 months ago

Resilience is not the same as high availability, so it can be in one AZ and still resilient (C)

Low latency (D)

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

upvoted 1 times

🗨️ 👤 **skywalker** 1 year, 11 months ago

Selected Answer: CD

CDCDCDCDCD

upvoted 4 times

🗨️ 👤 **skiwili** 2 years ago

Selected Answer: CD

C and D

upvoted 3 times

🗨️ 👤 **hiun** 2 years ago

Selected Answer: CD

C and D

upvoted 4 times

A company's customers are reporting increased latency while accessing static web content from Amazon S3. A SysOps administrator observed a very high rate of read operations on a particular S3 bucket.

What will minimize latency by reducing load on the S3 bucket?

- A. Migrate the S3 bucket to a region that is closer to end users' geographic locations.
- B. Use cross-region replication to replicate all of the data to another region.
- C. Create an Amazon CloudFront distribution with the S3 bucket as the origin.
- D. Use Amazon ElastiCache to cache data being served from Amazon S3.

Suggested Answer: C

Community vote distribution

C (100%)

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

Selected Answer: C

Clearly C

Rationale:

- A. Might help if everyone is always coming from "regions" that adjacent (requires effort/analysis)
 - B. Would increase load during replication, but might help later (same caveats as previous question)
 - C. Simplest/best because CF starts caching to every edge globally based only on actual user requests, bucket remains unchanged, and no analysis required.
 - D. ElastiCache is for DBs, not S3
- upvoted 7 times

 **jipark** Most Recent 10 months, 2 weeks ago

Selected Answer: C

why not A / B : question did not say "region" related traffic

- A. Migrate the S3 bucket to a region
- B. Use cross-region replication

upvoted 4 times

 **michaldavid** 1 year, 6 months ago

Selected Answer: C

cccccccc

upvoted 2 times

 **Pepepep** 1 year, 6 months ago

Discussion on this link.

<https://www.examttopics.com/discussions/amazon/view/7004-exam-aws-sysops-topic-1-question-611-discussion/>


upvoted 2 times

 **grka25** 1 year, 6 months ago

The answer is C. You can find this question in the old dump. Most people voted C there.

Cloudfront will cache the static content reducing the load on the S3 bucket

upvoted 2 times

 **beznika** 1 year, 6 months ago

C definitely. D doesn't even make any sense with static website. ElastiCache works together with a DataBase.



upvoted 1 times

 **CloudHandsOn** 1 year, 6 months ago

The answer should be D.

ElastiCache used with S3 for high reads is ideal to help with this kind of issue.

upvoted 1 times

  **Fatoch** 1 year, 7 months ago

Those questions are new .Where is the discussion ? after 120 no discussion. how do I believe default answer is correct. Anyone recently trying to get certificates

upvoted 4 times

  **BiggerStaff** 1 year, 7 months ago

I agree, why no discussion. How are the correct answers picked?

upvoted 1 times

A SysOps administrator needs to develop a solution that provides email notification and inserts a record into a database every time a file is put into an Amazon S3 bucket.



What is the MOST operationally efficient solution that meets these requirements?

- A. Set up an S3 event notification that targets an Amazon Simple Notification Service (Amazon SNS) topic. Create two subscriptions for the SNS topic. Use one subscription to send the email notification. Use the other subscription to invoke an AWS Lambda function that inserts the record into the database.
- B. Set up an Amazon CloudWatch alarm that enters ALARM state whenever an object is created in the S3 bucket. Configure the alarm to invoke an AWS Lambda function that sends the email notification and inserts the record into the database.
- C. Create an AWS Lambda function to send the email notification and insert the record into the database whenever a new object is detected in the S3 bucket. Invoke the function every minute with an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule.
- D. Set up two S3 event notifications. Target a separate AWS Lambda function with each notification. Configure one function to send the email notification. Configure the other function to insert the record into the database.

Suggested Answer: A

Community vote distribution



A (100%)

  **marcelodba** Highly Voted 2 years, 7 months ago

Selected Answer: A



<https://aws.amazon.com/pt/premiumsupport/knowledge-center/lambda-subscribe-sns-topic-same-account/>

upvoted 10 times

  **jipark** 1 year, 10 months ago

S3 itself support SNS



upvoted 1 times

  **10cc6ba** Most Recent 11 months, 3 weeks ago

Selected Answer: A

A only

upvoted 1 times

  **74d3443** 1 year, 6 months ago

A and D can be correct answer. IMO the term "Operationally Efficient" is a bit confusing in this scenario. I would go for A to manage one single notification from the source into a fan out strategy where I can handle the exception in SNS.

upvoted 1 times

  **74d3443** 1 year, 6 months ago

Please cancel this comment, I misread the D option. Is A for sure.

upvoted 1 times

  **edu_anadat** 2 years, 6 months ago

Selected Answer: A

You can't set up two S3 event notifications with the same Event types. But it's possible configure Lambda function, without SNS, in S3 event notifications

upvoted 4 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaaa

upvoted 2 times

A company hosts a web application on Amazon EC2 instances behind an Application Load Balancer. The instances are in an Amazon EC2 Auto Scaling group. The application is accessed with a public URL.

A SysOps administrator needs to implement a monitoring solution that checks the availability of the application and follows the same routes and actions as a customer. The SysOps administrator must receive a notification if less than 95% of the monitoring runs find no errors.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch Synthetics canary with a script that follows customer routes. Schedule the canary to run on a recurring schedule. Create a CloudWatch alarm that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic when the SuccessPercent metric is less than 95%.
- B. Create Amazon Route 53 health checks that monitor the availability of the endpoint. Create Amazon CloudWatch alarms that publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when the HealthCheckPercentageHealthy metric is less than 95%.
- C. Create a single AWS Lambda function to check whether the endpoints are available for each customer path. Schedule the Lambda function by using Amazon EventBridge (Amazon CloudWatch Events). Configure the Lambda function to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when an endpoint returns an error.
- D. Create an AWS Lambda function for each customer path to check whether that specific endpoint is available. Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events). Configure each Lambda function to publish a custom metric to Amazon CloudWatch for the endpoint status. Create CloudWatch alarms based on each custom metric to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when an alarm is in the ALARM state.

Suggested Answer: B

Community vote distribution

A (100%)

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

Selected Answer: A

A for sure

You can use Amazon CloudWatch Synthetics to create canaries, configurable scripts that run on a schedule, to monitor your endpoints and APIs. Canaries follow the same routes and perform the same actions as a customer, which makes it possible for you to continually verify your customer experience even when you don't have any customer traffic on your applications. By using canaries, you can discover issues before your customers do. https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Synthetics_Canaries.html

upvoted 13 times

 **jipark** Most Recent 10 months, 2 weeks ago

Selected Answer: A

Amazon CloudWatch Synthetics allows you to create canaries that simulate customer behavior and monitor the availability

why no B : health check monitor the availability of the endpoint. but cannot simulate customer's behavior.

upvoted 2 times

 **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaaa

upvoted 1 times

 **BugsBunny9998666** 1 year, 6 months ago

Selected Answer: A

canary for scripts

upvoted 1 times

 **hardwiredman** 1 year, 7 months ago

Selected Answer: A

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Synthetics_Canaries.html

upvoted 2 times

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

Think Canary in the coal mine...

upvoted 1 times

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

Selected Answer: A

Ans: A

Synthetics it's the only one who follows the same routes and actions as a customer.

upvoted 1 times

A SysOps administrator uses AWS Systems Manager Session Manager to connect to instances. After the SysOps administrator launches a new Amazon EC2 instance, the EC2 instance does not appear in the Session Manager list of systems that are available for connection. The SysOps administrator verifies that Systems Manager Agent is installed, updated, and running on the EC2 instance.

What is the reason for this issue?

- A. The SysOps administrator does not have access to the key pair that is required for connection.
- B. The SysOps administrator has not attached a security group to the EC2 instance to allow SSH on port 22.
- C. The EC2 instance does not have an attached IAM role that allows Session Manager to connect to the EC2 instance.
- D. The EC2 instance ID has not been entered into the Session Manager configuration.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Learning4life** 11 months ago

Selected Answer: C

Problem: After creating a new instance, the Session Manager tab in the Amazon Elastic Compute Cloud (Amazon EC2) console doesn't give you the option to connect.

Solution A: Create an instance profile: If you haven't already done so (as instructed by the information on the Session Manager tab in the EC2 console), create an AWS Identity and Access Management (IAM) instance profile by using Quick Setup. Quick Setup is a capability of AWS Systems Manager.

Session Manager requires an IAM instance profile to connect to your instance.

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

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: C

cccccc

upvoted 1 times

🗳️ 👤 **marcelodba** 2 years, 1 month ago

Selected Answer: C

It's c

upvoted 2 times

🗳️ 👤 **Jamshif01** 2 years, 1 month ago

Answer: C it is 100 %

upvoted 3 times

🗳️ 👤 **Liongeek** 2 years, 1 month ago

Ans: C!

upvoted 3 times

A SysOps administrator is unable to launch Amazon EC2 instances into a VPC because there are no available private IPv4 addresses in the VPC.

Which combination of actions must the SysOps administrator take to launch the instances? (Choose two.)

- A. Associate a secondary IPv4 CIDR block with the VPC.
- B. Associate a primary IPv6 CIDR block with the VPC.
- C. Create a new subnet for the VPC.
- D. Modify the CIDR block of the VPC.
- E. Modify the CIDR block of the subnet that is associated with the instances.

Suggested Answer: AC

Community vote distribution

AC (100%)

🗳️ 👤 **Liongeek** Highly Voted 1 year, 7 months ago

Ans: A&C

Every time you run out of IP's in your subnet, create a SECONDARY CIDR in your VPC. That will create a second local route in your RT's. Now you have more IP's available where you can create subnets.

Remember: you CANNOT modify the CIDR of a VPC or subnet.

Ref. <https://docs.aws.amazon.com/vpc/latest/userguide/configure-your-vpc.html#vpc-resize>

upvoted 16 times

🗳️ 👤 **alexiscloud** 8 months ago

absolutely

upvoted 1 times

🗳️ 👤 **noahsark** Most Recent 1 year, 4 months ago

Selected Answer: AC

<https://aws.amazon.com/premiumsupport/knowledge-center/subnet-insufficient-ips/>

upvoted 3 times

🗳️ 👤 **marcelodba** 1 year, 7 months ago

Selected Answer: AC

A, C looks right

upvoted 2 times

A SysOps administrator is creating an Amazon EC2 Auto Scaling group in a new AWS account. After adding some instances, the SysOps administrator notices that the group has not reached the minimum number of instances. The SysOps administrator receives the following error message:

Launching a new EC2 instance. Status Reason: Your quota allows for 0 more running instance(s).
You requested at least 1. Launching EC2 instance failed.

Which action will resolve this issue?

- A. Adjust the account spending limits for Amazon EC2 on the AWS Billing and Cost Management console.
- B. Modify the EC2 quota for that AWS Region in the EC2 Settings section of the EC2 console.
- C. Request a quota increase for the instance type family by using Service Quotas on the AWS Management Console.
- D. Use the Rebalance action in the Auto Scaling group on the AWS Management Console.

Suggested Answer: B

Community vote distribution

C (100%)

🗳️ 👤 **Liongeek** Highly Voted 1 year, 7 months ago

This one is tricky but I'll go for C.

1. You don't modify, you request an increase in your quota
2. Option B is too much general... there's no such EC2 Instance Quota per Region, the quota is per instance family.

upvoted 8 times

🗳️ 👤 **jipark** Most Recent 10 months, 2 weeks ago

Selected Answer: C

service limits(=quota) are preconfigured thresholds set by AWS to control resource usage by individual accounts

-> need AWS request

capacity limits refer to the available resources within an AWS service or region,

-> modify instance type, etc

upvoted 1 times

🗳️ 👤 **Andrew_A** 1 year ago

Selected Answer: C

The error message is suggesting that the account has hit its limit on the number of running EC2 instances. Therefore, the solution is to request a limit increase for the number of instances.

upvoted 1 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: C

cccccc

upvoted 1 times

🗳️ 👤 **goatbernard** 1 year, 6 months ago

Selected Answer: C

C i will choose

upvoted 1 times

🗳️ 👤 **marcelodba** 1 year, 7 months ago

Selected Answer: C


I'll go for C

upvoted 1 times

🗳️ 👤 **Fatoch** 1 year, 7 months ago

For me C too

upvoted 2 times

 **Xelnak** 1 year, 7 months ago

Selected Answer: C

C. Request a quota increase for the instance type family by using Service Quotas on the AWS Management Console.

upvoted 2 times

A SysOps administrator is creating two AWS CloudFormation templates. The first template will create a VPC with associated resources, such as subnets, route tables, and an internet gateway. The second template will deploy application resources within the VPC that was created by the first template. The second template should refer to the resources created by the first template.


How can this be accomplished with the LEAST amount of administrative effort?

- A. Add an export field to the outputs of the first template and import the values in the second template.
- B. Create a custom resource that queries the stack created by the first template and retrieves the required values.
- C. Create a mapping in the first template that is referenced by the second template.
- D. Input the names of resources in the first template and refer to those names in the second template as a parameter.

Suggested Answer: A


Community vote distribution

A (100%)

 **marcelodba** Highly Voted 2 years, 1 month ago

Selected Answer: A

A , You have to export first then import on 2nd template
upvoted 5 times

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

Selected Answer: A

Note: To reference a resource in another AWS CloudFormation stack, you must first create cross-stack references. To create a cross-stack reference, use the export field to flag the value of a resource output for export.

Ref link: <https://repost.aws/knowledge-center/cloudformation-reference-resource>

Ref link: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/walkthrough-crossstackref.html>

upvoted 5 times

 **james2033** Most Recent 10 months, 1 week ago

Selected Answer: A

'Export' output fields and 'Fn::ImportValue' , see <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/walkthrough-crossstackref.html>

(1) <https://s3.amazonaws.com/cloudformation-examples/user-guide/cross-stack/SampleNetworkCrossStack.template> --> 'Export'

(2) <https://s3.amazonaws.com/cloudformation-examples/user-guide/cross-stack/SampleWebAppCrossStack.template> --> 'Fn::ImportValue'

```
"Export": { "Name": { "Fn::Sub": "${AWS::StackName}-VPCID" } }
```

then

```
"GroupSet": [ { "Fn::ImportValue": { "Fn::Sub": "${NetworkStackName}-SecurityGroupID" } } ],
```

upvoted 3 times

 **noahsark** 1 year, 11 months ago

Selected Answer: A

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-reference-resource/>

upvoted 2 times

 **michaldavid** 2 years ago

Selected Answer: A

aaaaaa

upvoted 1 times

 **BugsBunny9998666** 2 years ago

Selected Answer: A

Its A, this question pop up in old topic for sys ops, there over 15 people voted A
upvoted 2 times

  **BugsBunny9998666** 2 years ago

here are some old comments

A is correct!



The optional Mappings section matches a key to a corresponding set of named values. For example, if you want to set values based on a region, you can create a mapping that uses the region name as a key and contains the values you want to specify for each specific region. You use the Fn::FindInMap intrinsic function to retrieve values in a map.

You can't include parameters, pseudo parameters, or intrinsic functions in the Mappings section.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/mappings-section-structure.html>

A is correct, mappings have nothing to do in referencing from one stack to the other

upvoted 1 times

  **Fatoch** 2 years, 1 month ago

It's c

Answer is C

upvoted 2 times

A company runs a web application on three Amazon EC2 instances behind an Application Load Balancer (ALB). The company notices that random periods of increased traffic cause a degradation in the application's performance. A SysOps administrator must scale the application to meet the increased traffic.

Which solution meets these requirements?

- A. Create an Amazon CloudWatch alarm to monitor application latency and increase the size of each EC2 instance if the desired threshold is reached.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor application latency and add an EC2 instance to the ALB if the desired threshold is reached.
- C. Deploy the application to an Auto Scaling group of EC2 instances with a target tracking scaling policy. Attach the ALB to the Auto Scaling group.
- D. Deploy the application to an Auto Scaling group of EC2 instances with a scheduled scaling policy. Attach the ALB to the Auto Scaling group.

Suggested Answer: C

Community vote distribution

C (95%)

5%

🗨️ **BugsBunny9998666** Highly Voted 2 years ago

Selected Answer: C

target tracking scaling policy is perfect for RANDOM periods of increased traffic.

For those who choose D, D is NOT an option here what are you ? Wizard with crystal ball to know when traffic will increase ?

upvoted 11 times

🗨️ **beznika** Highly Voted 2 years ago

It would be very useful to have downvote option as well.

upvoted 6 times

🗨️ **Christina666** Most Recent 1 year, 5 months ago

Selected Answer: C

random traffic-> target scheduling

upvoted 2 times

🗨️ **michaldavid** 2 years ago

Selected Answer: C

ccccccc

upvoted 1 times

🗨️ **beznika** 2 years ago

C is the right answer. How can it D if the option says scheduled scaling! How can you schedule scaling when you have random increase in traffic?

upvoted 2 times

🗨️ **marcelodba** 2 years, 1 month ago

Selected Answer: C

It's says random periods so schedule it's not an option

upvoted 4 times

🗨️ **Xelnak** 2 years, 1 month ago

Selected Answer: D

Scheduled scaling helps you to set up your own scaling schedule according to predictable load changes. For example, let's say that every week the traffic to your web application starts to increase on Wednesday, remains high on Thursday, and starts to decrease on Friday. You can configure a schedule for Amazon EC2 Auto Scaling to increase capacity on Wednesday and decrease capacity on Friday.

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

upvoted 1 times

🗨️ **AgboolaKun** 7 months, 1 week ago

You only use scheduled scaling only if traffic increase is predictable. The question says "the company notices that random periods of increased traffic", this shows that the time of traffic increase is unknown. Therefore, this is perfect use case for target scheduling.

upvoted 1 times

  **Liongeek** 2 years, 1 month ago

Ans: C

upvoted 2 times

  **Liongeek** 2 years, 1 month ago

My bad, Ans is D. It's says random periods so schedule it's not an option

upvoted 2 times

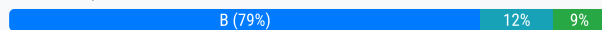
A company has a high-performance Windows workload. The workload requires a storage volume that provides consistent performance of 10,000 IOPS. The company does not want to pay for additional unneeded capacity to achieve this performance.

Which solution will meet these requirements with the LEAST cost?

- A. Use a Provisioned IOPS SSD (io1) Amazon Elastic Block Store (Amazon EBS) volume that is configured with 10,000 provisioned IOPS.
- B. Use a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume that is configured with 10,000 provisioned IOPS.
- C. Use an Amazon Elastic File System (Amazon EFS) file system in Max I/O mode.
- D. Use an Amazon FSx for Windows File Server file system that is configured with 10,000 IOPS.

Suggested Answer: D

Community vote distribution



JamesF92 Highly Voted 1 year, 4 months ago

Selected Answer: B

Use the official AWS cost calculator to compare. B costs way less than A or D for 10,000 IOPS. <https://calculator.aws/>
upvoted 5 times

stoy123 Most Recent 10 months, 1 week ago

Selected Answer: B

B
bbbbbb
upvoted 1 times

icecool36 10 months, 3 weeks ago

Selected Answer: B

Not D: They require a "storage volume" Fsx is not a storage volume.
Not A: This is much more expensive than B
NOT C: EFS is not for windows
B: With gp3 volumes, customers can scale IOPS (input/output operations per second) and throughput without needing to provision additional block storage capacity.
upvoted 3 times

Rabbit117 11 months ago

Selected Answer: B

General Purpose SSD (gp3) volumes are the latest generation of General Purpose SSD volumes, and the lowest cost SSD volume offered by Amazon EBS. This volume type helps to provide the right balance of price and performance for most applications. It also helps you to scale volume performance independently of volume size. This means that you can provision the required performance without needing to provision additional block storage capacity. Additionally, gp3 volumes offer a 20 percent lower price per GiB than General Purpose SSD (gp2) volumes.

gp3 volumes provide single-digit millisecond latency and 99.8 percent to 99.9 percent volume durability with an annual failure rate (AFR) no higher than 0.2 percent, which translates to a maximum of two volume failures per 1,000 running volumes over a one-year period. AWS designs gp3 volumes to deliver their provisioned performance 99 percent of the time.

upvoted 1 times

khaz123 11 months ago

Selected Answer: D

D with the description
upvoted 1 times

ordo 1 year, 1 month ago

B for sure
upvoted 1 times

jesusmoh 1 year, 2 months ago

For me is A.

Suitable for mission-critical workloads that require high performance and durability.

upvoted 1 times

  **Hatem08** 1 year, 1 month ago

but for least cost B makes better sense


upvoted 1 times

  **DocHolliday** 1 year, 4 months ago

Selected Answer: D

EBS are only for EC2 instances. The question does NOT specify any workload on a windows EC2...only that its a Windows workload....So would it not be D? I feel like for the question to be relevant to EBS they would HAVE to specify that its an EC2 instance else we have no way of knowing thats an option. But we know that D is possible.

upvoted 3 times

  **TwinSpark** 1 year, 2 months ago

I agree even the fact they say the do not want to pay for any unused capacity. Anycase very tricky question...

upvoted 1 times

  **Christina666** 1 year, 5 months ago

Selected Answer: B

gp3 is enough

upvoted 3 times

  **acethetest1000** 1 year, 5 months ago

Guys, why not D?

<https://docs.aws.amazon.com/fsx/latest/WindowsGuide/performance.html#performance-details-fsxw>



The company doesn't want to pay for unused space, hence they need an elastic file system.

upvoted 2 times

  **Student013657** 6 months, 3 weeks ago

Amazon EFS is not supported on Windows instances: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEFS.html>

upvoted 1 times

  **Gomer** 1 year, 8 months ago

Selected Answer: B



I'm no storage or math guru, but I did find links with formulas to calculate required minimum volume size to achieve 10,000 IOPS: 20 GiB for gp3 volume ($10,000 \text{ IOPS} / 500 \text{ IOPS per GiB} = 20 \text{ GiB}$) and 200 GiB for io1 volume ($10,000 \text{ IOPS} / 50 = 200 \text{ GiB}$).

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/general-purpose.html>

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/provisioned-iops.html>

If my math is wrong, correct me. However, since the question states that size matters, since 20 GiB is less than 200 GiB, I'd have to vote "B." gp3.

upvoted 3 times

  **csG13** 1 year, 9 months ago

Selected Answer: B

B - gp3 provides a consistent IOPS baseline performance that is not linearly dependent on the amount of storage we provision. It provides an initial 3000 IOPS baseline regardless and any IOPS above it costs 0.005\$.

As an example, assuming a size of 1TB and 10000 IOPS using a gp3 volume the monthly cost would be ~165\$. Using io1 the same would cost 825\$.

Evidently, the most cost effective solution is gp3.

upvoted 3 times

  **braveheart22** 1 year, 10 months ago

This is a tricky one, but I'm inclined toward AAAA

upvoted 1 times

  **Gil80** 1 year, 10 months ago

Selected Answer: A

A seems more suitable for the requirement

upvoted 1 times

  **noahsark** 1 year, 11 months ago

Selected Answer: B

Max IOPS per volume - 16,000

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html#vol-type-ssd>

upvoted 3 times

🗨️ 👤 **Brokdar** 1 year, 11 months ago

Selected Answer: A

Definitely A. gp3 does not offer consistent IOPS which is a requirement on the question. See below:

General Purpose SSD volumes (gp2 and gp3) balance price and performance for a wide variety of transactional workloads. These volumes are ideal for use cases such as boot volumes, medium-size single instance databases, and development and test environments.

Provisioned IOPS SSD volumes (io1 and io2) are designed to meet the needs of I/O-intensive workloads that are sensitive to storage performance and consistency. They provide a consistent IOPS rate that you specify when you create the volume.

upvoted 2 times

🗨️ 👤 **skywalker** 1 year, 11 months ago

Selected Answer: B

gp3 allow Max IOPS/Volume: 16,000

<https://aws.amazon.com/ebs/volume-types/>

upvoted 2 times

🗨️ 👤 **jipark** 1 year, 4 months ago

exactly !!

0~1600 is covered by General

upvoted 1 times

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

you need to pay to upgrade "Can scale up to 16,000 IOPS and 1,000 MiBps for an additional fee"

upvoted 1 times

A SysOps administrator must create a solution that automatically shuts down any Amazon EC2 instances that have less than 10% average CPU utilization for 60 minutes or more.

Which solution will meet this requirement in the MOST operationally efficient manner?

- A. Implement a cron job on each EC2 instance to run once every 60 minutes and calculate the current CPU utilization. Initiate an instance shutdown if CPU utilization is less than 10%.
- B. Implement an Amazon CloudWatch alarm for each EC2 instance to monitor average CPU utilization. Set the period at 1 hour, and set the threshold at 10%. Configure an EC2 action on the alarm to stop the instance.
- C. Install the unified Amazon CloudWatch agent on each EC2 instance, and enable the Basic level predefined metric set. Log CPU utilization every 60 minutes, and initiate an instance shutdown if CPU utilization is less than 10%.
- D. Use AWS Systems Manager Run Command to get CPU utilization from each EC2 instance every 60 minutes. Initiate an instance shutdown if CPU utilization is less than 10%.

Suggested Answer: B

Community vote distribution

B (100%)

🗲️ 👤 **zolthar_z** Highly Voted 👍 1 year ago

Selected Answer: B

Answer is B

upvoted 7 times

🗲️ 👤 **grka25** Most Recent 🕒 1 year ago

I will go with B. This is a simple example of AWS autoscaling.

upvoted 3 times

A SysOps administrator is unable to authenticate an AWS CLI call to an AWS service.

Which of the following is the cause of this issue?

- A. The IAM password is incorrect.
- B. The server certificate is missing.
- C. The SSH key pair is incorrect.
- D. There is no access key.

Suggested Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **KTsankov** Highly Voted 👍 11 months, 1 week ago

Selected Answer: D

When the AWS CLI runs a command, it sends an encrypted request to the AWS servers to perform the appropriate AWS service operations. Your credentials (the ACCESS key and secret key) are involved in the encryption and enable AWS to authenticate the person making the request. There are several things that can interfere with the correct operation of this process, as follows.

<https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-troubleshooting.html>

upvoted 6 times

🗲️ 👤 **Liongeek** Highly Voted 👍 1 year, 7 months ago

Ans: D

You don't use passwords in AWS CLI. You use access and secret access key.

upvoted 5 times

🗲️ 👤 **Christina666** Most Recent 🕒 11 months, 1 week ago

Selected Answer: D

~/.aws config file stores key

upvoted 3 times

🗲️ 👤 **dehkon** 1 year, 2 months ago

Signature Version 4 is the process to add authentication information to AWS requests sent by HTTP. For security, most requests to AWS must be signed with an access key, which consists of an access key ID and secret access key. These two keys are commonly referred to as your security credentials. For information about how to obtain credentials for your account, see Understanding and getting your credentials.

upvoted 2 times

🗲️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

dddddd

upvoted 3 times

A company requires that all IAM user accounts that have not been used for 90 days or more must have their access keys and passwords immediately disabled. A SysOps administrator must automate the process of disabling unused keys using the MOST operationally efficient method.

How should the SysOps administrator implement this solution?

- A. Create an AWS Step Functions workflow to identify IAM users that have not been active for 90 days. Run an AWS Lambda function when a scheduled Amazon EventBridge (Amazon CloudWatch Events) rule is invoked to automatically remove the AWS access keys and passwords for these IAM users.
- B. Configure an AWS Config rule to identify IAM users that have not been active for 90 days. Set up an automatic weekly batch process on an Amazon EC2 instance to disable the AWS access keys and passwords for these IAM users.
- C. Develop and run a Python script on an Amazon EC2 instance to programmatically identify IAM users that have not been active for 90 days. Automatically delete these IAM users.
- D. Set up an AWS Config managed rule to identify IAM users that have not been active for 90 days. Set up an AWS Systems Manager automation runbook to disable the AWS access keys for these IAM users.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **ede83d8** 4 months, 3 weeks ago

Selected Answer: D

AWSConfigRemediation-DeleteIAMUser
upvoted 1 times

🗳️ **Christina666** 11 months, 1 week ago

Selected Answer: D

Checks if your AWS Identity and Access Management (IAM) users have passwords or active access keys that have not been used within the specified number of days you provided. The rule is NON_COMPLIANT if there are inactive accounts not recently used.
upvoted 1 times

🗳️ **Tony183** 11 months, 2 weeks ago

Selected Answer: D

DDDDDDDD
upvoted 2 times

🗳️ **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddd
upvoted 2 times

🗳️ **Pepepep** 1 year, 6 months ago

D.
<https://docs.aws.amazon.com/config/latest/developerguide/iam-user-unused-credentials-check.html>
upvoted 4 times

🗳️ **marcelodba** 1 year, 7 months ago

Selected Answer: D

I'll go for D
upvoted 3 times

A company creates custom AMI images by launching new Amazon EC2 instances from an AWS CloudFormation template. It installs and configures necessary software through AWS OpsWorks, and takes images of each EC2 instance. The process of installing and configuring software can take between 2 to 3 hours, but at times, the process stalls due to installation errors.

The SysOps administrator must modify the CloudFormation template so if the process stalls, the entire stack will fail and roll back.


Based on these requirements, what should be added to the template?

- A. Conditions with a timeout set to 4 hours.
- B. CreationPolicy with a timeout set to 4 hours.
- C. DependsOn with a timeout set to 4 hours.
- D. Metadata with a timeout set to 4 hours.

Suggested Answer: B

Community vote distribution

B (100%)

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

Selected Answer: B

When you provision an Amazon EC2 instance in an AWS CloudFormation stack, you might specify additional actions to configure the instance, such as install software packages or bootstrap applications. Normally, CloudFormation proceeds with stack creation after the instance has been successfully created. However, you can use a CreationPolicy so that CloudFormation proceeds with stack creation only after your configuration actions are done. That way you'll know your applications are ready to go after stack creation succeeds.

A CreationPolicy instructs CloudFormation to wait on an instance until CloudFormation receives the specified number of signals. This policy takes effect only when CloudFormation creates the instance.

Ref link: <https://aws.amazon.com/blogs/devops/use-a-creationpolicy-to-wait-for-on-instance-configurations/>

Ref link: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-creationpolicy.html>

upvoted 8 times

 **tamng** Most Recent 12 months ago

B is correct

upvoted 1 times

 **michaldavid** 2 years ago

Selected Answer: B

bbbbbbb

upvoted 2 times

 **Raynor** 2 years, 1 month ago

Selected Answer: B

B - https://docs.aws.amazon.com/ko_kr/AWSCloudFormation/latest/UserGuide/aws-attribute-creationpolicy.html

upvoted 4 times

A company runs workloads on 90 Amazon EC2 instances in the eu-west-1 Region in an AWS account. In 2 months, the company will migrate the workloads from eu-west-1 to the eu-west-3 Region.

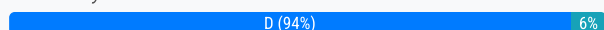
The company needs to reduce the cost of the EC2 instances. The company is willing to make a 1-year commitment that will begin next week. The company must choose an EC2 instance purchasing option that will provide discounts for the 90 EC2 instances regardless of Region during the 1-year period.

Which solution will meet these requirements?

- A. Purchase EC2 Standard Reserved Instances.
- B. Purchase an EC2 Instance Savings Plan.
- C. Purchase EC2 Convertible Reserved Instances.
- D. Purchase a Compute Savings Plan.

Suggested Answer: C

Community vote distribution



Christina666 Highly Voted 1 year, 11 months ago

Selected Answer: D

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, size, AZ, Region, operating system, or tenancy, and also apply to Fargate and Lambda usage. For example, with Compute Savings Plans, you can change from C4 to M5 instances, shift a workload from EU (Ireland) to Europe (London), or move a workload from Amazon EC2 to Fargate or Lambda at any time and automatically continue to pay the Savings Plans price.

upvoted 7 times

braveheart22 Highly Voted 2 years, 4 months ago

D is the right option

Compute saving plan is regardless of the region.

upvoted 6 times

10cc6ba Most Recent 11 months, 3 weeks ago

Selected Answer: B

B: Purchase an EC2 Instance Savings Plan

Coverage Across Regions: EC2 Instance Savings Plans provide flexibility in terms of instance type and Region, allowing coverage for instances running in any AWS Region.

Commitment Term: EC2 Instance Savings Plans offer a 1-year term, which aligns with the company's requirement for a 1-year commitment.

Savings Benefit: Savings Plans provide discounts based on your usage, regardless of instance family, size, AZ, or Region. This means as long as the instances are running within the scope of the Savings Plan, they will benefit from the discount, even when migrated from eu-west-1 to eu-west-3.

upvoted 1 times

Aamee 8 months ago

you haven't mentioned how much interms of percent of savings option B can provide compared to the option of Compute Savings plan in option D?... plus, no source link provided for your justification here either...

upvoted 1 times

hiun 2 years, 5 months ago

Selected Answer: D

D. Global

upvoted 2 times

tts1234 2 years, 6 months ago

Selected Answer: D

D. Purchase a Compute Savings Plan.



upvoted 2 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddd

upvoted 2 times


  **Xelnak** 2 years, 7 months ago

Selected Answer: D

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).

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

upvoted 4 times

  **Liongeek** 2 years, 7 months ago

Ans: D

Compute Saving Plans are global

The other ones are regional.

Question asks for global

upvoted 5 times

A SysOps administrator has created a VPC that contains a public subnet and a private subnet. Amazon EC2 instances that were launched in the private subnet cannot access the internet. The default network ACL is active on all subnets in the VPC, and all security groups allow all outbound traffic.

Which solution will provide the EC2 instances in the private subnet with access to the internet?

- A. Create a NAT gateway in the public subnet. Create a route from the private subnet to the NAT gateway.
- B. Create a NAT gateway in the public subnet. Create a route from the public subnet to the NAT gateway.
- C. Create a NAT gateway in the private subnet. Create a route from the public subnet to the NAT gateway.
- D. Create a NAT gateway in the private subnet. Create a route from the private subnet to the NAT gateway.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Liongeek** Highly Voted 👍 2 years, 7 months ago

Ans: A

upvoted 8 times

🗳️ 👤 **Aamee** Most Recent ⌚ 8 months ago

Selected Answer: A

100% option A is correct here..

upvoted 1 times

🗳️ 👤 **10cc6ba** 11 months, 3 weeks ago

Selected Answer: A

A only as answer

upvoted 1 times

🗳️ 👤 **strovertz** 1 year, 7 months ago

Selected Answer: A

ANS is A

upvoted 2 times

🗳️ 👤 **jipark** 1 year, 10 months ago

Selected Answer: A

NAT gateway should place on public subnet.

source is on private subnet.

connect from private to NAT gateway.

upvoted 4 times

🗳️ 👤 **Tony183** 1 year, 11 months ago

Selected Answer: A

A is correct

upvoted 3 times

🗳️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaa

upvoted 3 times

A company plans to run a public web application on Amazon EC2 instances behind an Elastic Load Balancer (ELB). The company's security team wants to protect the website by using AWS Certificate Manager (ACM) certificates. The ELB must automatically redirect any HTTP requests to HTTPS.

Which solution will meet these requirements?

- A. Create an Application Load Balancer that has one HTTPS listener on port 80. Attach an SSL/TLS certificate to listener port 80. Create a rule to redirect requests from HTTP to HTTPS.
- B. Create an Application Load Balancer that has one HTTP listener on port 80 and one HTTPS protocol listener on port 443. Attach an SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.
- C. Create an Application Load Balancer that has two TCP listeners on port 80 and port 443. Attach an SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.
- D. Create a Network Load Balancer that has two TCP listeners on port 80 and port 443. Attach an SSL/TLS certificate to listener port 443. Create a rule to redirect requests from port 80 to port 443.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **Rhydian25** 8 months ago

Selected Answer: B

It is B as the HTTP protocol uses the port 80 and the HTTPS uses the 443

upvoted 1 times

🗳️ 👤 **xSohox** 1 year, 4 months ago

Selected Answer: B

That is B.

upvoted 4 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: B

bbbbbbbbb

upvoted 3 times

🗳️ 👤 **marcelodba** 2 years, 1 month ago

Selected Answer: B

I'll go for B

upvoted 4 times

A company wants to track its AWS costs in all member accounts that are part of an organization in AWS Organizations. Managers of the member accounts want to receive a notification when the estimated costs exceed a predetermined amount each month. The managers are unable to configure a billing alarm. The IAM permissions for all users are correct.

What could be the cause of this issue?


- A. The management/payer account does not have billing alerts turned on.
- B. The company has not configured AWS Resource Access Manager (AWS RAM) to share billing information between the member accounts and the management/payer account.
- C. Amazon GuardDuty is turned on for all the accounts.
- D. The company has not configured an AWS Config rule to monitor billing.

Suggested Answer: A

Community vote distribution

A (90%)


10%

 **KTsankov** Highly Voted 11 months, 1 week ago

Selected Answer: A

In a consolidated billing account, member linked account metrics are captured only if the payer account enables the Receive Billing Alerts preference. If you change which account is your management/payer account, you must enable the billing alerts in the new management/payer account.

upvoted 5 times

 **Christina666** Most Recent 11 months, 1 week ago

Selected Answer: A

Before you can create an alarm for your estimated charges, you must enable billing alerts, so that you can monitor your estimated AWS charges and create an alarm.

After you enable billing alerts for the first time, it takes about 15 minutes before you can view billing data and set billing alarms.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/monitor_estimated_charges_with_cloudwatch.html#turning_on_billing_metrics:~:text=

upvoted 4 times


 **bakamon** 11 months, 3 weeks ago

Selected Answer: B

B is the answer. By default, the management/payer account has access to billing information for all the accounts in the organization. It can enable billing alerts and set up notifications for cost thresholds.

However, in this scenario, the managers of the member accounts are unable to configure a billing alarm. This indicates that the issue lies with the permissions and access of the managers, rather than the configuration of the management/payer account.

upvoted 1 times

 **eboehm** 11 months, 2 weeks ago

what the heck are you talking about? for starters it states that the permissions for all users are correct. Secondly, what does that have to do with Resource Access Manager(RAM), which is used for securely shared your provisioned aws resources and NOT billing information!

upvoted 9 times

 **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaaaaa

upvoted 1 times

 **marcelodba** 1 year, 7 months ago

Selected Answer: A

A is correct

upvoted 1 times

🔖 👤 **Liongeek** 1 year, 7 months ago

Ans: A

Ref.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/monitor_estimated_charges_with_cloudwatch.html#turning_on_billing_metrics

upvoted 3 times

A company is using Amazon Elastic Container Service (Amazon ECS) to run a containerized application on Amazon EC2 instances. A SysOps administrator needs to monitor only traffic flows between the ECS tasks.

Which combination of steps should the SysOps administrator take to meet this requirement? (Choose two.)

- A. Configure Amazon CloudWatch Logs on the elastic network interface of each task.
- B. Configure VPC Flow Logs on the elastic network interface of each task.
- C. Specify the awsvpc network mode in the task definition.
- D. Specify the bridge network mode in the task definition.
- E. Specify the host network mode in the task definition.

Suggested Answer: AC

Community vote distribution

BC (100%)

🗳️ 👤 **Arnaud92** Highly Voted 📌 2 years ago

B,C : <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-networking-awsvpc.html>

upvoted 5 times

🗳️ 👤 **Rabbit117** Most Recent 🕒 9 months, 3 weeks ago

Selected Answer: BC

The task networking features that are provided by the awsvpc network mode give Amazon ECS tasks the same networking properties as Amazon EC2 instances. Using the awsvpc network mode simplifies container networking, because you have more control over how your applications communicate with each other and other services within your VPCs. The awsvpc network mode also provides greater security for your containers by allowing you to use security groups and network monitoring tools at a more granular level within your tasks. You can also use other Amazon EC2 networking features such as VPC Flow Logs to monitor traffic to and from your tasks.

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task-networking-awsvpc.html>

upvoted 1 times

🗳️ 👤 **jipark** 1 year, 4 months ago

Selected Answer: BC

B. VPC Flow Logs capture IP traffic going to and from network interfaces in your VPC

C. The awsvpc network mode allows each task to have its own dedicated network namespace

upvoted 2 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: BC

The awsvpc network mode also provides greater security for your containers by enabling you to use security groups and network monitoring tools at a more granular level within your tasks. Because each task gets its own elastic network interface (ENI), you can also use other Amazon EC2 networking features such as VPC Flow Logs to monitor traffic to and from your tasks. Additionally, containers that belong to the same task can communicate over the localhost interface.

upvoted 1 times

🗳️ 👤 **Gomer** 1 year, 8 months ago

Selected Answer: BC

"The awsvpc network mode also provides greater security for your containers by enabling you to use security groups and network monitoring tools at a more granular level within your tasks. Because each task gets its own elastic network interface (ENI), you can also use other Amazon EC2 networking features such as VPC Flow Logs to monitor traffic to and from your tasks."

upvoted 2 times

🗳️ 👤 **hiun** 2 years ago

Selected Answer: BC

B, C is correct answer

upvoted 2 times

🗳️ 👤 **marcelodba** 2 years, 1 month ago

Selected Answer: BC

I'll go for B,C

upvoted 3 times

A company uses AWS Organizations to manage multiple AWS accounts. The company's SysOps team has been using a manual process to create and manage IAM roles. The team requires an automated solution to create and manage the necessary IAM roles for multiple AWS accounts.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create AWS CloudFormation templates. Reuse the templates to create the necessary IAM roles in each of the AWS accounts.
- B. Use AWS Directory Service with AWS Organizations to automatically associate the necessary IAM roles with Microsoft Active Directory users.
- C. Use AWS Resource Access Manager with AWS Organizations to deploy and manage shared resources across the AWS accounts.
- D. Use AWS CloudFormation StackSets with AWS Organizations to deploy and manage IAM roles for the AWS accounts.

Suggested Answer: D


Community vote distribution

D (100%)

  **jipark** Highly Voted 10 months, 2 weeks ago

Selected Answer: D

Stackset for Organization.
upvoted 6 times

  **jipark** 10 months, 2 weeks ago

StackSets can be targeted to specific organizational units (OUs) within AWS Organizations
upvoted 2 times

  **michaldavid** Most Recent 1 year, 6 months ago

Selected Answer: D

ddddddd
upvoted 2 times

  **marcelodba** 1 year, 7 months ago

Selected Answer: D

I'll go for D
<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-concepts.html>
upvoted 3 times

A SysOps administrator needs to configure automatic rotation for Amazon RDS database credentials. The credentials must rotate every 30 days. The solution must integrate with Amazon RDS.

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

- A. Store the credentials in AWS Systems Manager Parameter Store as a secure string. Configure automatic rotation with a rotation interval of 30 days.
- B. Store the credentials in AWS Secrets Manager. Configure automatic rotation with a rotation interval of 30 days.
- C. Store the credentials in a file in an Amazon S3 bucket. Deploy an AWS Lambda function to automatically rotate the credentials every 30 days.
- D. Store the credentials in AWS Secrets Manager. Deploy an AWS Lambda function to automatically rotate the credentials every 30 days.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ **tttfakil** 10 months, 3 weeks ago

Selected Answer: B

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-secrets-manager.html>

upvoted 2 times

🗳️ **Christina666** 11 months, 1 week ago

Selected Answer: B

Secret manager has build-in feature of automately rotate password

upvoted 4 times

🗳️ **CVDON** 1 year, 1 month ago

D because rds only works with Secret Manager

upvoted 1 times

🗳️ **CVDON** 1 year, 1 month ago

Only Works With secret manager

upvoted 1 times

🗳️ **King_AWS_202** 1 year, 1 month ago

Selected Answer: B

The answer is definitely B

upvoted 3 times

🗳️ **Gomer** 1 year, 1 month ago

For what it's worth, Lambda is involved with rotation of secrets (suggesting option "D."), but I'm presuming Lambda is configured transparently in the background when you enable key rotation, therefore "B." seems to be a more straightforward answer.

"Secrets Manager rotation uses an AWS Lambda function to update the secret and the database."

"To rotate a secret, Secrets Manager calls a Lambda function according to the schedule you set up."

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

upvoted 1 times

🗳️ **marylynn** 1 year, 2 months ago

Ans is D

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

upvoted 1 times

🗳️ **michaldavid** 1 year, 6 months ago

Selected Answer: B

bbbbbbb

upvoted 1 times

🗳️ **marcelodba** 1 year, 7 months ago

Selected Answer: B

Ans. B

upvoted 1 times

A company's SysOps administrator attempts to restore an Amazon Elastic Block Store (Amazon EBS) snapshot. However, the snapshot is missing because another system administrator accidentally deleted the snapshot. The company needs the ability to recover snapshots for a specified period of time after snapshots are deleted.

Which solution will provide this functionality?

- A. Turn on deletion protection on individual EBS snapshots that need to be kept.
- B. Create an IAM policy that denies the deletion of EBS snapshots by using a condition statement for the snapshot age. Apply the policy to all users.
- C. Create a Recycle Bin retention rule for EBS snapshots for the desired retention period.
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to copy EBS snapshots to Amazon S3 Glacier.

Suggested Answer: C

Community vote distribution

C (100%)

🗲️ 👤 **Christina666** Highly Voted 👍 11 months, 1 week ago

Selected Answer: C

New Recycle Bin

In order to give you more control over the deletion process, we are launching a Recycle Bin for EBS Snapshots. As you will see in a moment, you can now set up rules to retain deleted snapshots so that you can recover them after an accidental deletion. You can think of this as a two-level model, where individual AWS users are responsible for the initial deletion, and then a designated "Recycle Bin Administrator" (as specified by an IAM role) manages retention and recovery.

upvoted 6 times

🗲️ 👤 **michaldavid** Most Recent 🕒 1 year, 6 months ago

Selected Answer: C

cccccccc

upvoted 1 times

🗲️ 👤 **marcelodba** 1 year, 7 months ago

Selected Answer: C

<https://aws.amazon.com/pt/blogs/aws/new-recycle-bin-for-ebs-snapshots/>

upvoted 2 times

🗲️ 👤 **Liongeek** 1 year, 7 months ago

Ans: C

REF. <https://aws.amazon.com/es/blogs/aws/new-recycle-bin-for-ebs-snapshots/>

upvoted 3 times

A SysOps administrator recently configured Amazon S3 Cross-Region Replication on an S3 bucket.

Which of the following does this feature replicate to the destination S3 bucket by default?

- A. Objects in the source S3 bucket for which the bucket owner does not have permissions
- B. Objects that are stored in S3 Glacier
- C. Objects that existed before replication was configured
- D. Object metadata

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Christina666** Highly Voted 👍 11 months, 1 week ago

Selected Answer: D

By default, Amazon S3 replicates the following:

Objects created after you add a replication configuration.

Unencrypted objects.

Objects encrypted using customer provided keys (SSE-C), objects encrypted at rest under an Amazon S3 managed key (SSE-S3) or a KMS key stored in AWS Key Management Service (SSE-KMS). For more information, see Replicating objects created with server-side encryption (SSE-C, SSE-S3, SSE-KMS, DSSE-KMS).

Object metadata from the source objects to the replicas. For information about replicating metadata from the replicas to the source objects, see Replicating metadata changes with Amazon S3 replica modification sync.

Only objects in the source bucket for which the bucket owner has permissions to read objects and access control lists (ACLs).
upvoted 7 times

🗳️ 👤 **Christina666** Most Recent ⌚ 11 months, 1 week ago

Selected Answer: D

Object ACL updates, unless you direct Amazon S3 to change the replica ownership when source and destination buckets aren't owned by the same accounts.

For more information, see Changing the replica owner.

It can take a while until Amazon S3 can bring the two ACLs in sync. This change in ownership applies only to objects created after you add a replication configuration to the bucket.

Object tags, if there are any.

S3 Object Lock retention information, if there is any.
upvoted 2 times

🗳️ 👤 **Gomer** 1 year, 2 months ago

Selected Answer: D

What is replicated by default?

"Object metadata from the source objects to the replicas."


"Only objects in the source bucket for which the bucket owner has permissions to read objects and access control lists (ACLs)."
upvoted 2 times

🗳️ 👤 **defmania00** 1 year, 4 months ago

Selected Answer: D

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-what-is-isnot-replicated.html#replication-what-is-replicated>

upvoted 3 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddd

upvoted 1 times



  **gsotiriou** 1 year, 6 months ago

I think D as well

[https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-what-is-isnot-](https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-what-is-isnot-replicated.html#:~:text=Object%20metadata%20from%20the%20source%20objects%20to%20the%20replicas.%20For%20information%20about%20replicating)

[replicated.html#:~:text=Object%20metadata%20from%20the%20source%20objects%20to%20the%20replicas.%20For%20information%20about%20replicating](https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-what-is-isnot-replicated.html#:~:text=Object%20metadata%20from%20the%20source%20objects%20to%20the%20replicas.%20For%20information%20about%20replicating)

upvoted 1 times

  **Fatoch** 1 year, 6 months ago



In this case it is not A. By default S3 bucket does not have permissions. Why D? I didn't get yet

upvoted 1 times

  **marylynn** 1 year, 2 months ago

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-what-is-isnot-replicated.html#replication-what-is-replicated>

upvoted 1 times

  **grka25** 1 year, 6 months ago

I think D.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html>

upvoted 1 times

A company has a workload that is sending log data to Amazon CloudWatch Logs. One of the fields includes a measure of application latency. A SysOps administrator needs to monitor the p90 statistic of this field over time.

What should the SysOps administrator do to meet this requirement?

- A. Create an Amazon CloudWatch Contributor Insights rule on the log data.
- B. Create a metric filter on the log data.
- C. Create a subscription filter on the log data.
- D. Create an Amazon CloudWatch Application Insights rule for the workload.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **Christina666** Highly Voted 👍 11 months, 1 week ago

Selected Answer: B

Statistics are metric data aggregations over specified periods of time. When you graph or retrieve the statistics for a metric, you specify the Period of time, such as five minutes, to use to calculate each statistical value. For example, if the Period is five minutes, the Sum is the sum of all sample values collected during the five-minute period, while the Minimum is the lowest value collected during the five-minute period.

Percentile (p) indicates the relative standing of a value in a dataset. For example, p95 is the 95th percentile and means that 95 percent of the data within the period is lower than this value and 5 percent of the data is higher than this value. Percentiles help you get a better understanding of the distribution of your metric data.

upvoted 6 times

🗳️ 👤 **noahsark** Highly Voted 👍 1 year, 3 months ago

Selected Answer: B

Create a metric filter on the log data.

Percentile (p) indicates the relative standing of a value in a dataset. For example, p95 is the 95th percentile and means that 95 percent of the data within the period is lower than this value and 5 percent of the data is higher than this value. Percentiles help you get a better understanding of the distribution of your metric data.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Statistics-definitions.html>

upvoted 6 times

🗳️ 👤 **jipark** Most Recent 🕒 10 months, 2 weeks ago

Selected Answer: B

metric filter is enough

upvoted 1 times

🗳️ 👤 **Pepepep** 1 year, 6 months ago

B

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html>

upvoted 1 times

A company wants to archive sensitive data on Amazon S3 Glacier. The company's regulatory and compliance requirements do not allow any modifications to the data by any account.

Which solution meets these requirements?

- A. Attach a vault lock policy to an S3 Glacier vault that contains the archived data. Use the lock ID to validate the vault lock policy after 24 hours.
- B. Attach a vault lock policy to an S3 Glacier vault that contains the archived data. Use the lock ID to validate the vault lock policy within 24 hours.
- C. Configure S3 Object Lock in governance mode. Upload all files after 24 hours.
- D. Configure S3 Object Lock in governance mode. Upload all files within 24 hours.

Suggested Answer: A

Community vote distribution

B (100%)


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

Selected Answer: B

While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires. To prevent your vault from exiting the in-progress state, you must complete the Vault Lock process within these 24 hours. Otherwise, your Vault Lock policy will be deleted.

<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock.html#vault-lock-overview>

upvoted 10 times

 **Christina666** Most Recent 11 months, 1 week ago


Selected Answer: B

Locking a vault takes two steps:

Initiate the lock by attaching a Vault Lock policy to your vault, which sets the lock to an in-progress state and returns a lock ID. While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires. To prevent your vault from exiting the in-progress state, you must complete the Vault Lock process within these 24 hours. Otherwise, your Vault Lock policy will be deleted.

Use the lock ID to complete the lock process. If the Vault Lock policy doesn't work as expected, you can stop the Vault Lock process and restart from the beginning. For information about how to use the S3 Glacier API to lock a vault, see Locking a Vault by Using the S3 Glacier API.

upvoted 3 times

 **Gomer** 1 year, 2 months ago

Selected Answer: B

Only Glacier Vault Lock Policy can block any user from deleting a file regardless of age or other circumstance.

S3 Object lock: "With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary."

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html>

S3 Glacier Vault Lock Access policy:

"Vault access policy that can be locked. After you lock a Vault Lock policy, the policy can't be changed. You can use a Vault Lock Policy to enforce compliance controls."

"Locking a vault takes two steps:"

1. "attaching a Vault Lock policy to your vault, which"... "returns a lock ID"... "you must complete the Vault Lock process within these 24 hours."
2. "Use the lock ID to complete the lock process."

<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock.html>

upvoted 3 times

 **noahsark** 1 year, 3 months ago

Selected Answer: B

Attach a vault lock policy to an S3 Glacier vault that contains the archived data. Use the lock ID to validate the vault lock policy within 24 hours.

Notes:

While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires.

<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock.html>

upvoted 2 times

🗨️ 👤 **Fatoch** 1 year, 6 months ago

A and B are same answers.

upvoted 2 times

🗨️ 👤 **Gomer** 1 year, 2 months ago

No they are not. Before and after 24 hours makes all the difference.

upvoted 1 times

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

Selected Answer: B

<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock.html#:~:text=Initiate%20the%20lock,will%20be%20deleted.>

upvoted 1 times

🗨️ 👤 **Beidog** 1 year, 7 months ago

Selected Answer: B

Vote for B

upvoted 1 times

🗨️ 👤 **Raynor** 1 year, 7 months ago

B - <https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock.html#vault-lock-overview>

upvoted 2 times

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

Ans: A

Ref: <https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock.html#vault-lock-overview>

upvoted 1 times

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

Based on your link the answer is B

upvoted 2 times

A company manages an application that uses Amazon ElastiCache for Redis with two extra-large nodes spread across two different Availability Zones. The company's IT team discovers that the ElastiCache for Redis cluster has 75% freeable memory. The application must maintain high availability.

What is the MOST cost-effective way to resize the cluster?

- A. Decrease the number of nodes in the ElastiCache for Redis cluster from 2 to 1.
- B. Deploy a new ElastiCache for Redis cluster that uses large node types. Migrate the data from the original cluster to the new cluster. After the process is complete, shut down the original cluster.
- C. Deploy a new ElastiCache for Redis cluster that uses large node types. Take a backup from the original cluster, and restore the backup in the new cluster. After the process is complete, shut down the original cluster.
- D. Perform an online resizing for the ElastiCache for Redis cluster. Change the node types from extra-large nodes to large nodes.

Suggested Answer: A

Community vote distribution

D (100%)

🗳️ 👤 **Xelnak** Highly Voted 1 year, 7 months ago

Selected Answer: D

https://docs.amazonaws.cn/en_us/AmazonElastiCache/latest/red-ug/redis-cluster-vertical-scaling-scaling-down.html#redis-cluster-vertical-scaling-down-console

upvoted 9 times

🗳️ 👤 **Christina666** Most Recent 11 months, 1 week ago

Selected Answer: D

Scaling down Redis cache clusters (Console)

Scaling down Redis cache clusters (Amazon CLI)

Scaling down Redis cache clusters (ElastiCache API)

upvoted 4 times

🗳️ 👤 **edu_anadat** 1 year, 6 months ago

Selected Answer: D

D

The application must maintain high availability.

upvoted 4 times

🗳️ 👤 **Beidog** 1 year, 7 months ago

Selected Answer: D

Answer is D

upvoted 2 times

A company must migrate its applications to AWS. The company is using Chef recipes for configuration management. The company wants to continue to use the existing Chef recipes after the applications are migrated to AWS.

What is the MOST operationally efficient solution that meets these requirements?

- A. Use AWS CloudFormation to create an Amazon EC2 instance, install a Chef server, and add Chef recipes.
- B. Use AWS CloudFormation to create a stack and add layers for Chef recipes.
- C. Use AWS Elastic Beanstalk with the Docker platform to upload Chef recipes.
- D. Use AWS OpsWorks to create a stack and add layers with Chef recipes.

Suggested Answer: D

Community vote distribution

D (100%)

fig 7 months, 2 weeks ago

Selected Answer: D

This is probably an old/obsolete question now...

AWS OpsWorks for Chef Automate is no longer accepting new customers. Existing customers will be unaffected until May 5, 2024 at which time the service will become unavailable. We recommend that existing customers migrate to Chef SaaS or an alternative solution.

Source:

https://docs.aws.amazon.com/opsworks/latest/userguide/welcome_opscm.html

upvoted 2 times

dangji 1 year, 5 months ago

Selected Answer: D

The key word is "Chef".

upvoted 3 times

michaldavid 1 year, 6 months ago

Selected Answer: D

ddddddd

upvoted 1 times

marcelodba 1 year, 7 months ago

Selected Answer: D

https://docs.aws.amazon.com/opsworks/latest/userguide/welcome_opscm.html

upvoted 1 times

A company uses AWS Organizations to manage its AWS accounts. A SysOps administrator must create a backup strategy for all Amazon EC2 instances across all the company's AWS accounts.

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

- A. Deploy an AWS Lambda function to each account to run EC2 instance snapshots on a scheduled basis.
- B. Create an AWS CloudFormation stack set in the management account to add an AutoBackup=True tag to every EC2 instance.
- C. Use AWS Backup in the management account to deploy policies for all accounts and resources.
- D. Use a service control policy (SCP) to run EC2 instance snapshots on a scheduled basis in each account.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Christina666** 11 months, 1 week ago

Selected Answer: C

Backup role-based access control

With AWS Backup, a backup operator can back up all supported resources on AWS without requiring the backup operator to have direct access to those resources. This provides a separation of control where resource owners can't impact the retention of backups, and backup operators can't mutate or exfiltrate data.

You can set resource-based access policies on backup vaults. With resource-based access policies, you can control access to backups in a backup vault across all users, rather than having to define permissions for each user.

upvoted 3 times

🗳️ 👤 **Gomer** 1 year, 2 months ago

Selected Answer: C

"You can delegate backup policy management in AWS Organizations and cross account monitoring in AWS Backup. This enables delegating backup management to a dedicated backup administration account, removing the need for member accounts to access management accounts for backup administration. Delegated backup administrators can create and manage backup policies and monitor backup activity across accounts. Organization-wide backup administration delegation through AWS Organizations enables securely centralized backup management at scale."

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

upvoted 1 times

🗳️ 👤 **Xelnak** 1 year, 7 months ago

Selected Answer: C

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

AWS Backup now supports cross-account backup, enabling you to securely copy your backups across your AWS accounts within your AWS organizations.

upvoted 4 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

Ans C

Tagging isn't enough to backup the instances. You'll also need a backup plan and rules.

upvoted 1 times

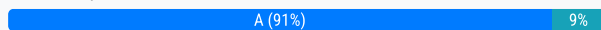
A SysOps administrator is reviewing VPC Flow Logs to troubleshoot connectivity issues in a VPC. While reviewing the logs, the SysOps administrator notices that rejected traffic is not listed.

What should the SysOps administrator do to ensure that all traffic is logged?

- A. Create a new flow log that has a filter setting to capture all traffic.
- B. Create a new flow log. Set the log record format to a custom format. Select the proper fields to include in the log.
- C. Edit the existing flow log. Change the filter setting to capture all traffic.
- D. Edit the existing flow log. Set the log record format to a custom format. Select the proper fields to include in the log.

Suggested Answer: C

Community vote distribution



🗳️ 👤 **Liongeek** Highly Voted 2 years, 1 month ago

Ans: A

You CANNOT modify a VPC Flow Log

Ref: <https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html#flow-logs-limitations> "After you create a flow log, you cannot change its configuration or the flow log record format. For example, you can't associate a different IAM role with the flow log, or add or remove fields in the flow log record. Instead, you can delete the flow log and create a new one with the required configuration."

upvoted 18 times

🗳️ 👤 **jipark** 1 year, 4 months ago

cleared !!

"You CANNOT modify a VPC Flow Log"

upvoted 1 times

🗳️ 👤 **Rabbit117** Most Recent 9 months, 3 weeks ago

Selected Answer: A

I think the answer is A.

Flow logs basics:

You can create a flow log for a VPC, a subnet, or a network interface. If you create a flow log for a subnet or VPC, each network interface in that subnet or VPC is monitored.

Flow log data for a monitored network interface is recorded as flow log records, which are log events consisting of fields that describe the traffic flow. For more information, see Flow log records.

To create a flow log, you specify:

The resource for which to create the flow log.

The type of traffic to capture (accepted traffic, rejected traffic, or all traffic).

The destinations to which you want to publish the flow log data.

<https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html>

upvoted 1 times

🗳️ 👤 **shmulik** 1 year, 3 months ago

Selected Answer: A

After you create a flow log, you cannot change its configuration or the flow log record format. For example, you can't associate a different IAM role with the flow log, or add or remove fields in the flow log record. Instead, you can delete the flow log and create a new one with the required configuration.

upvoted 1 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: A

Flow logs can help you with a number of tasks, such as:

Diagnosing overly restrictive security group rules

Monitoring the traffic that is reaching your instance

Determining the direction of the traffic to and from the network interfaces

Flow log data is collected outside of the path of your network traffic, and therefore does not affect network throughput or latency. You can create or delete flow logs without any risk of impact to network performance.

upvoted 2 times

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

Selected Answer: A

Answer: A

upvoted 1 times

🗨️ 👤 **csG13** 1 year, 8 months ago

Selected Answer: A

A - you can't modify an existing flow log. Also it's not B since we want to log all traffic.

upvoted 1 times

🗨️ 👤 **gulu73** 1 year, 11 months ago

Selected Answer: A

Ans A

You cannot modify VPC Flow Log just tested it.

upvoted 1 times

🗨️ 👤 **CodePoet** 1 year, 12 months ago

Selected Answer: A

Keyword: "ensure that all traffic is logged"

upvoted 1 times

🗨️ 👤 **yeacuz** 2 years ago

Selected Answer: B

The answer is:

B. Create a new flow log. Set the log record format to a custom format. Select the proper fields to include in the log.

You need to create a new flow log. There is no "filter setting to capture all traffic", but you can set the log record to a custom format as specified in the user guide (<https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html#flow-logs-custom>)

upvoted 1 times

🗨️ 👤 **csG13** 1 year, 9 months ago

There is; there are three types of traffic - accepted traffic, rejected traffic, or all traffic. Correct answer is A.

See more here: <https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html>

upvoted 2 times

🗨️ 👤 **marcelodba** 2 years ago

Selected Answer: A

Ans: A

upvoted 2 times

A company is expanding its use of AWS services across its portfolios. The company wants to provision AWS accounts for each team to ensure a separation of business processes for security, compliance, and billing. Account creation and bootstrapping should be completed in a scalable and efficient way so new accounts are created with a defined baseline and governance guardrails in place. A SysOps administrator needs to design a provisioning process that saves time and resources.

Which action should be taken to meet these requirements?

- A. Automate using AWS Elastic Beanstalk to provision the AWS accounts, set up infrastructure, and integrate with AWS Organizations.
- B. Create bootstrapping scripts in AWS OpsWorks and combine them with AWS CloudFormation templates to provision accounts and infrastructure.
- C. Use AWS Config to provision accounts and deploy instances using AWS Service Catalog.
- D. Use AWS Control Tower to create a template in Account Factory and use the template to provision new accounts.

Suggested Answer: D

Community vote distribution

D (100%)

 **Christina666** 11 months, 1 week ago

Selected Answer: D

If you are hosting more than a handful of accounts, it's beneficial to have an orchestration layer that facilitates account deployment and account governance. You can adopt AWS Control Tower as your primary way to provision accounts and infrastructure. With AWS Control Tower, you can more easily adhere to corporate standards, meet regulatory requirements, and follow best practices.

AWS Control Tower enables end users on your distributed teams to provision new AWS accounts quickly, by means of configurable account templates in Account Factory. Meanwhile, your central cloud administrators can monitor that all accounts are aligned with established, company-wide compliance policies.

upvoted 4 times

 **noahsark** 1 year, 3 months ago

Selected Answer: D

Use AWS Control Tower to create a template in Account Factory and use the template to provision new accounts.

https://d1.awsstatic.com/products/control-tower/Product-Page-Diagram_AWS-Control-Tower.9281926228bb2900c76b4b6d85b2819efc078978.png

upvoted 2 times

 **jipark** 10 months, 2 weeks ago

great !!

upvoted 1 times

 **braveheart22** 1 year, 4 months ago

The answer is DDD.

Account Factory – An Account Factory is a configurable account template that helps to standardize the provisioning of new accounts with pre-approved account configurations. AWS Control Tower offers a built-in Account Factory that helps automate the account provisioning workflow in your organization.

<https://docs.aws.amazon.com/controltower/latest/userguide/what-is-control-tower.html>


upvoted 1 times

 **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddd


upvoted 1 times

 **tyfta6** 1 year, 6 months ago

Selected Answer: D

Elastic Beanstalk is for Deploying Web Apps, NOT for automating MultiAccount AWS Infrastructures (which is all that Control Tower does)

upvoted 2 times

  **Liongeek** 1 year, 7 months ago

Ans: D

upvoted 1 times

A SysOps administrator noticed that the cache hit ratio for an Amazon CloudFront distribution is less than 10%.

Which collection of configuration changes will increase the cache hit ratio for the distribution? (Choose two.)

- A. Ensure that only required cookies, query strings, and headers are forwarded in the Cache Behavior Settings.
- B. Change the Viewer Protocol Policy to use HTTPS only.
- C. Configure the distribution to use presigned cookies and URLs to restrict access to the distribution.
- D. Enable automatic compression of objects in the Cache Behavior Settings.
- E. Increase the CloudFront time to live (TTL) settings in the Cache Behavior Settings.

Suggested Answer: AE

Community vote distribution

AE (100%)

🗳️ 👤 **Christina666** 11 months, 1 week ago

Selected Answer: AE

A cache hit refers to the situation wherein the cache is able to successfully retrieve data and content that was saved to it, and then display it on a web page

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cache-hit-ratio.html>

Specifying how long CloudFront caches your objects

Using Origin Shield

Caching based on query string parameters

Caching based on cookie values

Caching based on request headers

Remove Accept-Encoding header when compression is not needed

Serving media content by using HTTP

upvoted 3 times

🗳️ 👤 **noahsark** 1 year, 3 months ago

Selected Answer: AE

Increase the CloudFront time to live (TTL) settings in the Cache Behavior Settings.

Ensure that only required cookies, query strings, and headers are forwarded in the Cache Behavior Settings.

By default, each file automatically expires after 24 hours

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

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.html#DownloadDistValuesCacheBehavior>

upvoted 1 times

🗳️ 👤 **csG13** 1 year, 4 months ago

Selected Answer: AE

A and E

upvoted 1 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: AE

A and E

upvoted 1 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

Ans: A&E

upvoted 3 times

A SysOps administrator is attempting to download patches from the internet into an instance in a private subnet. An internet gateway exists for the VPC, and a NAT gateway has been deployed on the public subnet; however, the instance has no internet connectivity. The resources deployed into the private subnet must be inaccessible directly from the public internet.

Public Subnet (10.0.1.0/24) Route Table

Destination Target -

10.0.0.0/16 local

0.0.0.0/0 IGW

Private Subnet (10.0.2.0/24) Route Table

Destination Target -

10.0.0.0/16 local

What should be added to the private subnet's route table in order to address this issue, given the information provided?

- A. 0.0.0.0/0 IGW
- B. 0.0.0.0/0 NAT
- C. 10.0.1.0/24 IGW
- D. 10.0.1.0/24 NAT

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ **tgiv** 8 months, 1 week ago

Selected Answer: B

The correct answer is 0.0.0.0/0 NAT.

upvoted 1 times

🗳️ **Christina666** 1 year, 5 months ago

Selected Answer: B

To enable instances in a private subnet to connect to the internet, you can create a NAT gateway or launch a NAT instance in a public subnet. Then add a route for the private subnet's route table that routes IPv4 internet traffic (0.0.0.0/0) to the NAT device.

<https://docs.aws.amazon.com/vpc/latest/userguide/route-table-options.html#route-tables-nat>

upvoted 4 times

🗳️ **Christina666** 1 year, 5 months ago

why 0.0.0.0/0?

upvoted 2 times

🗳️ **k0s3k** 1 year, 5 months ago

default route == 0.0.0.0/0 == internet

So you say, that the internet traffic should be routed through NAT

upvoted 3 times

🗳️ **braveheart22** 1 year, 10 months ago

Oh yeah the answer is BBB

upvoted 1 times

🗳️ **michaldavid** 2 years ago

Selected Answer: B

bbbbbbbbb

upvoted 2 times

🗨️ 👤 **Jamshif01** 2 years, 1 month ago

B is the correct answer

private ip --> nat gateway --> 0.0.0.0/0 --> internet access

upvoted 3 times

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

Ans: A

upvoted 1 times

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

My Bad, correct Ans is B

upvoted 3 times

A company is undergoing an external audit of its systems, which run wholly on AWS. A SysOps administrator must supply documentation of Payment Card Industry Data Security Standard (PCI DSS) compliance for the infrastructure managed by AWS.

Which set of actions should the SysOps administrator take to meet this requirement?

- A. Download the applicable reports from the AWS Artifact portal and supply these to the auditors.
- B. Download complete copies of the AWS CloudTrail log files and supply these to the auditors.
- C. Download complete copies of the AWS CloudWatch logs and supply these to the auditors.
- D. Provide the auditors with administrative access to the production AWS account so that the auditors can determine compliance.

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Liongeek** Highly Voted 🏆 1 year, 7 months ago

Ans: A

Just notice that this is correct as is regarding AWS infrastructure, but if you want to know if your resource and account configuration is compliant, you should then check AWS Security Hub

upvoted 8 times

🗨️ 👤 **Christina666** Most Recent 🔍 11 months, 1 week ago

Selected Answer: A

The PCI DSS Attestation of Compliance (AOC) and Responsibility Summary is available to customers through AWS Artifact, a self-service portal for on-demand access to AWS compliance reports. Sign in to AWS Artifact in the AWS Management Console, or learn more at Getting Started with AWS Artifact.

upvoted 3 times

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

Selected Answer: A

Ans= A

upvoted 2 times

🗨️ 👤 **Pacoca** 1 year, 4 months ago

Answer= A

<https://aws.amazon.com/compliance/pci-dss-level-1-faqs/>

upvoted 1 times

A company has an initiative to reduce costs associated with Amazon EC2 and AWS Lambda.

Which action should a SysOps administrator take to meet these requirements?

- A. Analyze the AWS Cost and Usage Report by using Amazon Athena to identify cost savings.
- B. Create an AWS Budgets alert to alarm when account spend reaches 80% of the budget.
- C. Purchase Reserved Instances through the Amazon EC2 console.
- D. Use AWS Compute Optimizer and take action on the provided recommendations.

Suggested Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **Liongeek** Highly Voted 👍 1 year, 7 months ago

Ans: D

Compute Optimizer gives recommendations for EC2, EBS, Lambda

upvoted 8 times

🗲️ 👤 **jipark** Most Recent ⌚ 10 months, 2 weeks ago

Selected Answer: D

why not A : athena do not give recommendation.

upvoted 4 times

🗲️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddd

upvoted 4 times

A company wants to use only IPv6 for all its Amazon EC2 instances. The EC2 instances must not be accessible from the internet, but the EC2 instances must be able to access the internet. The company creates a dual-stack VPC and IPv6-only subnets.

How should a SysOps administrator configure the VPC to meet these requirements?

- A. Create and attach a NAT gateway. Create a custom route table that includes an entry to point all IPv6 traffic to the NAT gateway. Attach the custom route table to the IPv6-only subnets.
- B. Create and attach an internet gateway. Create a custom route table that includes an entry to point all IPv6 traffic to the internet gateway. Attach the custom route table to the IPv6-only subnets.
- C. Create and attach an egress-only internet gateway. Create a custom route table that includes an entry to point all IPv6 traffic to the egress-only internet gateway. Attach the custom route table to the IPv6-only subnets.
- D. Create and attach an internet gateway and a NAT gateway. Create a custom route table that includes an entry to point all IPv6 traffic to the internet gateway and all IPv4 traffic to the NAT gateway. Attach the custom route table to the IPv6-only subnets.

Suggested Answer: D

Community vote distribution

C (100%)

🗳️ **Rabbit117** 9 months, 3 weeks ago

Selected Answer: C

IPv6 addresses are public by default therefore they are accessible from the internet. If you want an IPv6 instance to have access to the internet but not be accessible from the internet you should use an IPv6 egress-only internet gateway. Then add a route in the route table to point internet traffic, ::/0 to the egress-only internet gateway.

<https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

upvoted 2 times

🗳️ **jipark** 1 year, 4 months ago

Selected Answer: C

IPv6 only (egress) : Internet G/W

IPv4 only (ingress) : NAT G/w

upvoted 3 times

🗳️ **Christina666** 1 year, 5 months ago

Selected Answer: C

Egress-only internet gateway basics

IPv6 addresses are globally unique, and are therefore public by default. If you want your instance to be able to access the internet, but you want to prevent resources on the internet from initiating communication with your instance, you can use an egress-only internet gateway. To do this, create an egress-only internet gateway in your VPC, and then add a route to your route table that points all IPv6 traffic (::/0) or a specific range of IPv6 address to the egress-only internet gateway. IPv6 traffic in the subnet that's associated with the route table is routed to the egress-only internet gateway.

An egress-only internet gateway is stateful: it forwards traffic from the instances in the subnet to the internet or other AWS services, and then sends the response back to the instances.

upvoted 4 times

🗳️ **fazlur21** 1 year, 6 months ago

Selected Answer: C

Answer is C

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

upvoted 2 times

🗳️ **Andrew_A** 1 year, 6 months ago

Selected Answer: C

NAT Gateways are primarily used for IPv4 traffic and not IPv6

upvoted 1 times

🗳️ **AndyMartinez** 1 year, 8 months ago

Selected Answer: C

C. egress-only internet gateway
upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: C

cccccc
upvoted 1 times

  **BugsBunny9998666** 2 years ago

Selected Answer: C

ingress iP v4 only
egress iP v6 only /// one way connection to internet without being exposed
upvoted 1 times

  **XAvenger** 2 years ago

Selected Answer: C

egress-only internet gateway
upvoted 1 times

  **hardwiredman** 2 years, 1 month ago

Selected Answer: C

IPV6 needs an internet-only gateway
upvoted 3 times

  **Liongeek** 2 years, 1 month ago

Ans is C
Only egress-only internet gateway can be used to let instance go to internet without being exposed
upvoted 4 times

A company has an existing web application that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB) across two Availability Zones. The application uses an Amazon RDS Multi-AZ DB Instance. Amazon Route 53 record sets route requests for dynamic content to the load balancer and requests for static content to an Amazon S3 bucket. Site visitors are reporting extremely long loading times.

Which actions should be taken to improve the performance of the website? (Choose two.)

- A. Add Amazon CloudFront caching for static content.
- B. Change the load balancer listener from HTTPS to TCP.
- C. Enable Amazon Route 53 latency-based routing.
- D. Implement Amazon EC2 Auto Scaling for the web servers.
- E. Move the static content from Amazon S3 to the web servers.

Suggested Answer: AD

Community vote distribution

AD (90%)

10%

🗳️ 👤 **zolthar_z** Highly Voted 1 year, 11 months ago

Selected Answer: AD

AD:

A is obvious

D If servers are working to total capacity the ASG will help to increase the performance,

If is a latency issue C will not work because both are in the same region

upvoted 9 times

🗳️ 👤 **tgiv** 8 months, 1 week ago

I agree, makes sense.

upvoted 1 times

🗳️ 👤 **BugsBunny9998666** Highly Voted 2 years ago

A D /// C is wrong as it mentioned that (ALB) across two Availability Zones) JUST 1 region !!!

what is the point in C, Route 53 latency routing if it all goes to one region anyway ?

upvoted 6 times

🗳️ 👤 **jipark** Most Recent 1 year, 4 months ago

Selected Answer: AD

why not C : latency based route to fast responsive "Region".

but multi-AZ is one region.

upvoted 3 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: AD

static: cloudfront

dynamic: increase ASG app performance

upvoted 3 times

🗳️ 👤 **noahsark** 1 year, 9 months ago

Selected Answer: AD

Add Amazon CloudFront caching for static content.

Implement Amazon EC2 Auto Scaling for the web servers.

Not Enable Amazon Route 53 latency-based routing.:

"To use latency-based routing, you create latency records for your resources in multiple AWS Regions."

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-latency.html>

upvoted 2 times

🗨️ **michaldavid** 2 years ago

Selected Answer: AD

A and D

upvoted 1 times

🗨️ **beznika** 2 years ago

AD 100%

upvoted 1 times

🗨️ **tyfta6** 2 years, 1 month ago

Selected Answer: AD

True A. Add Amazon CloudFront caching for static content. (For S3)

Wrong B. Change the load balancer listener from HTTPS to TCP. (ALB not supported TCP. NLB supported TCP and has extreme performance)

Wrong C. Enable Amazon Route 53 latency-based routing. (Application is in one region. Dont need latency)

True D. Implement Amazon EC2 Auto Scaling for the web servers. (Auto Scaling can control app performance by scale out and scale in)

Wrong E. Move the static content from Amazon S3 to the web servers.

upvoted 3 times

🗨️ **marcelodba** 2 years, 1 month ago

Selected Answer: AC

I'll go for A,C

upvoted 2 times

🗨️ **Jamshif01** 2 years, 1 month ago

ans: AC

A. Add Amazon CloudFront caching for static content. -YES

B. Change the load balancer listener from HTTPS to TCP. - doesn't make any sense

C. Enable Amazon Route 53 latency-based routing. - YES

D. Implement Amazon EC2 Auto Scaling for the web servers. - This won't help with loading

E. Move the static content from Amazon S3 to the web servers. - no, cloudfront caching going to solve this problem

upvoted 2 times

🗨️ **Liongeek** 2 years, 1 month ago

This on is tricky. I think it's A & C, but not sure if it's also D.

upvoted 2 times

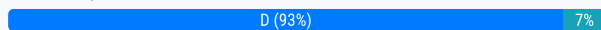
A company is running an application on premises and wants to use AWS for data backup. All of the data must be available locally. The backup application can write only to block-based storage that is compatible with the Portable Operating System Interface (POSIX).

Which backup solution will meet these requirements?

- A. Configure the backup software to use Amazon S3 as the target for the data backups.
- B. Configure the backup software to use Amazon S3 Glacier as the target for the data backups.
- C. Use AWS Storage Gateway, and configure it to use gateway-cached volumes.
- D. Use AWS Storage Gateway, and configure it to use gateway-stored volumes.

Suggested Answer: D

Community vote distribution



skywalker Highly Voted 1 year, 5 months ago

Selected Answer: D

Stored volumes make your entire data available locally on the gateway, while maintaining an asynchronous copy in the S3 bucket. Cached volumes store the full volume in the S3 bucket, while only keeping the recently used data in local cache.

upvoted 11 times

csG13 Highly Voted 1 year, 4 months ago

Selected Answer: D

The answer is D.

A, B are clearly not suitable. In terms of C vs D, the questions states

"All of the data must be available locally"

When a volume gateway works in cached mode it only stores the frequently accessed data locally for low-latency access, not the entire dataset. On the other hand, with stored volume gateway you backup all your data locally.

upvoted 5 times

Christina666 Most Recent 11 months, 1 week ago

Selected Answer: D

data available on local-> stored volume

upvoted 2 times

edgardopm 1 year, 5 months ago

Selected Answer: D

Volume Gateway provides an iSCSI target, which enables you to create block storage volumes and mount them as iSCSI devices from your on-premises or EC2 application servers. The Volume Gateway runs in either a cached or stored mode.

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.

upvoted 5 times

MrMLB 1 year, 6 months ago

Selected Answer: C

C. Use AWS Storage Gateway, and configure it to use gateway-cached volumes.

AWS Storage Gateway is a service that connects an on-premises data center to the cloud. It provides block-based storage that is compatible with the Portable Operating System Interface (POSIX) and can be used as a target for data backups. Gateway-cached volumes allow you to store your data locally, while asynchronously backing up that data to Amazon S3. This allows you to retain local access to your data, while still providing the benefits of cloud-based data storage and backup.

upvoted 2 times

🗋️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

dddddd

upvoted 1 times

🗋️ 👤 **Xelnak** 1 year, 7 months ago

Selected Answer: D

https://docs.amazonaws.cn/en_us/storagegateway/latest/vgw/StorageGatewayConcepts.html

By using stored volumes, you can store your primary data locally, while asynchronously backing up that data to Amazon. Stored volumes provide your on-premises applications with low-latency access to their entire datasets. At the same time, they provide durable, offsite backups.

upvoted 3 times

🗋️ 👤 **Liongeek** 1 year, 7 months ago

Ans: D

With Volume Stored Gateway all your data stays onprem and backups are sent to S3 as EBS snapshots. It's also POSIX compatible

upvoted 3 times

A global company handles a large amount of personally identifiable information (PII) through an internal web portal. The company's application runs in a corporate data center that is connected to AWS through an AWS Direct Connect connection. The application stores the PII in Amazon S3. According to a compliance requirement, traffic from the web portal to Amazon S3 must not travel across the internet.

What should a SysOps administrator do to meet the compliance requirement?

- A. Provision an interface VPC endpoint for Amazon S3. Modify the application to use the interface endpoint.
- B. Configure AWS Network Firewall to redirect traffic to the internal S3 address.
- C. Modify the application to use the S3 path-style endpoint.
- D. Set up a range of VPC network ACLs to redirect traffic to the internal S3 address.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Xelnak** Highly Voted 1 year, 7 months ago

Selected Answer: A

Using the interface endpoint, applications in your on-premises data center can easily query S3 buckets over AWS Direct Connect or Site-to-Site VPN.
<https://aws.amazon.com/blogs/architecture/choosing-your-vpc-endpoint-strategy-for-amazon-s3/>
upvoted 6 times

🗳️ 👤 **Christina666** Highly Voted 11 months, 1 week ago

Selected Answer: A

Key words:

data traffic not cross internet-> S3 interface endpoint
If inside VPC, then S3 gateway endpoint
upvoted 6 times

🗳️ 👤 **satamex** 9 months, 1 week ago

Your explanations are always great.. kudos..
upvoted 4 times

🗳️ 👤 **michaldavid** Most Recent 1 year, 6 months ago

Selected Answer: A

aaaaaa
upvoted 2 times

🗳️ 👤 **beznika** 1 year, 6 months ago

This question is written wrong. For S3 there is no interface VPC endpoint. S3 and DynamoDB uses VPC Gateway Endpoint. Interface VPC endpoints require ENI and S3 doesn't use it.
upvoted 2 times

🗳️ 👤 **beznika** 1 year, 6 months ago

A answer is still correct but I was wrong about the S3, it can use both VPC gateway endpoint and interface VPC endpoint.
upvoted 3 times

🗳️ 👤 **Liongeek** 1 year, 7 months ago

Ans: A!
upvoted 3 times

A SysOps administrator notices a scale-up event for an Amazon EC2 Auto Scaling group. Amazon CloudWatch shows a spike in the RequestCount metric for the associated Application Load Balancer. The administrator would like to know the IP addresses for the source of the requests.

Where can the administrator find this information?

- A. Auto Scaling logs
- B. AWS CloudTrail logs
- C. EC2 instance logs
- D. Elastic Load Balancer access logs

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Liongeek** Highly Voted 1 year, 7 months ago

Ans: D!

upvoted 5 times

🗳️ 👤 **fazlur21** Highly Voted 1 year ago

Answer D

"Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and troubleshoot issues."

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

upvoted 5 times

🗳️ 👤 **28b8844** Most Recent 3 months, 1 week ago

Selected Answer: C

The ELB Log needs to be enabled explicitly. The question does not mention that the ELB access logs are enabled at this moment.

upvoted 1 times

🗳️ 👤 **alexiscloud** 7 months, 4 weeks ago

Answer: D

upvoted 1 times

🗳️ 👤 **Christina666** 11 months, 1 week ago

Selected Answer: D

Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and troubleshoot issues.

upvoted 4 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: D

dddddd

upvoted 2 times

A company's SysOps administrator deploys a public Network Load Balancer (NLB) in front of the company's web application. The web application does not use any Elastic IP addresses. Users must access the web application by using the company's domain name. The SysOps administrator needs to configure Amazon Route 53 to route traffic to the NLB.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a Route 53 AAAA record for the NLB.
- B. Create a Route 53 alias record for the NLB.
- C. Create a Route 53 CAA record for the NLB.
- D. Create a Route 53 CNAME record for the NLB.

Suggested Answer: B

Community vote distribution

B (100%)

  **BugsBunny9998666** Highly Voted 1 year, 6 months ago

Selected Answer: B

As a good reminder:

A record = URL to IPv4

AAAA record = URL to IPv6

CNAME record = URL to URL (All the same, one url = Many URL's)



Alias record = AWS service

upvoted 39 times

  **DennisRichard** Most Recent 8 months, 2 weeks ago

In the question, I don't understand the emphasis on cost though

upvoted 2 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: B

bbbbbb

upvoted 2 times

  **Liongeek** 1 year, 7 months ago

Ans: B

Alias record doesn't charge when you associate it to AWS resources like NLB

upvoted 2 times

A company runs an encrypted Amazon RDS for Oracle DB instance. The company wants to make regular backups available in another AWS Region.

What is the MOST operationally efficient solution that meets these requirements?

- A. Modify the DB instance. Enable cross-Region automated backups.
- B. Create an RDS read replica in another Region. Create a snapshot of the read replica.
- C. Use AWS Database Migration Service (AWS DMS) to copy the data to a DB instance in another Region.
- D. Temporarily turn off encryption on the DB instance. Take a snapshot. Copy the snapshot to another Region.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Liongeek** Highly Voted 1 year, 7 months ago

Ans: A

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

"Enabling cross-Region automated backups

You can enable backup replication on new or existing DB instances using the Amazon RDS console. You can also use the start-db-instance-automated-backups-replication AWS CLI command or the StartDBInstanceAutomatedBackupsReplication RDS API operation. You can replicate up to 20 backups to each destination AWS Region for each AWS account"

upvoted 11 times

🗳️ 👤 **JamesF92** Most Recent 10 months, 2 weeks ago

Selected Answer: A

the thing is ... question is asking for "backups" in another region. A read replica is for active use by db clients' live read operations. We don't want that ... we just want a backup.

upvoted 4 times

🗳️ 👤 **Christina666** 11 months ago

Selected Answer: A

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReplicateBackups.html#AutomatedBackups.Replicating.Enable

upvoted 1 times

🗳️ 👤 **awsguru1998** 1 year, 4 months ago

In terms of operational efficiency, Option A may be a better solution compared to Option B. However, it's important to note that cross-Region backups may cause additional latency, and the backups may take longer to complete if there is limited bandwidth between the Regions. The company should assess its specific needs and consider the trade-offs between automation and latency when selecting the best solution for its requirements.

upvoted 3 times

🗳️ 👤 **ahalamri** 1 year, 2 months ago

Do replicas work cross-region ? I believe its expensive not cost effective

upvoted 1 times

🗳️ 👤 **jas26says** 1 year ago

For what I understood in the question, cost was never something to consider in this case.

upvoted 2 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: A

aaaaaaa

upvoted 2 times

A company is rolling out a new version of its website. Management wants to deploy the new website in a limited rollout to 20% of the company's customers. The company uses Amazon Route 53 for its website's DNS solution.


Which configuration will meet these requirements?

- A. Create a failover routing policy. Within the policy, configure 80% of the website traffic to be sent to the original resource. Configure the remaining 20% of traffic as the failover record that points to the new resource.
- B. Create a multivalue answer routing policy. Within the policy, create 4 records with the name and IP address of the original resource. Configure 1 record with the name and IP address of the new resource.
- C. Create a latency-based routing policy. Within the policy, configure a record pointing to the original resource with a weight of 80. Configure a record pointing to the new resource with a weight of 20.
- D. Create a weighted routing policy. Within the policy, configure a weight of 80 for the record pointing to the original resource. Configure a weight of 20 for the record pointing to the new resource.

Suggested Answer: D

Community vote distribution

D (100%)

  **alexiscloud** 7 months, 4 weeks ago

Answer: D

upvoted 1 times

  **michaldavid** 1 year, 6 months ago

Selected Answer: D

ddddddd

upvoted 2 times

  **marcelodba** 1 year, 7 months ago

Selected Answer: D

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-weighted.html>

upvoted 2 times

  **Liongeek** 1 year, 7 months ago

Ans: D

upvoted 4 times

A SysOps administrator created an AWS CloudFormation template that provisions Amazon EC2 instances, an Elastic Load Balancer (ELB), and an Amazon RDS DB instance. During stack creation, the creation of the EC2 instances and the creation of the ELB are successful. However, the creation of the DB instance fails.

What is the default behavior of CloudFormation in this scenario?

- A. CloudFormation will roll back the stack and delete the stack.
- B. CloudFormation will roll back the stack but will not delete the stack.
- C. CloudFormation will prompt the user to roll back the stack or continue.
- D. CloudFormation will successfully complete the stack but will report a failed status for the DB instance.

Suggested Answer: B

Community vote distribution

B (76%)

A (24%)

  **traja**  2 years, 7 months ago

Selected Answer: B

Ans is B. It will delete all the resources in the stack but stack deletion will be manual.

upvoted 10 times

  **dspd** 2 months, 3 weeks ago

tested in my account

upvoted 1 times

  **Vivec**  2 years, 3 months ago

Selected Answer: B

By default, when a resource creation fails during the creation of a CloudFormation stack, CloudFormation will roll back the stack to its previous state, undoing any changes made to the stack.

Therefore, the answer is B. CloudFormation will roll back the stack but will not delete the stack. This allows the administrator to troubleshoot the issue that caused the failure and attempt to create the stack again.

upvoted 5 times

  **Rado_Piatek**  1 month, 1 week ago

Selected Answer: B

Ans: B

Resources are ofc deleted by default

Tested it plus i had to rewatch video for it to make sure, essentially u can review logs, which means stack is NOT deleted by default if creation fails.

What is important, after creation failure, stacks number WILL increment, what means stack itself doesn't get deleted.

BTW same goes for updating the stack, you still can review events and logs, it will be rolled back to latest working state.



upvoted 1 times

  **dspd** 2 months, 3 weeks ago

Selected Answer: A

A is the answer

upvoted 1 times

  **numark** 7 months, 1 week ago

Answer is A and ChatGPT agrees: By default, when a CloudFormation stack creation fails for any resource, CloudFormation rolls back the entire stack. This means that CloudFormation will attempt to undo all the changes it has made so far and will delete any resources that were successfully created (like the EC2 instances and the ELB in your scenario). The stack itself will also be deleted.

upvoted 1 times

  **ivaldesgo** 8 months ago

Selected Answer: A

CloudFormation by default delete everything after the rollback.

upvoted 1 times

🗨️ 👤 **ivaldesgo** 8 months ago

CloudFormation by default delete everything after the rollback.

upvoted 2 times

🗨️ 👤 **Aamee** 9 months, 2 weeks ago

Selected Answer: B

B for sure!

upvoted 1 times

🗨️ 👤 **10cc6ba** 11 months, 3 weeks ago

Selected Answer: A

A. CloudFormation will roll back the stack and delete the stack.

Explanation:

When a stack creation fails, AWS CloudFormation automatically rolls back all resources that were created up to the point of failure. This behavior is intended to ensure that partial and potentially unstable stacks are not left in your account.

This rollback operation includes deleting any resources that were successfully created before the failure occurred. This means that both the EC2 instances and the Elastic Load Balancer will be deleted since the creation of the DB instance failed.

This automatic rollback is the default behavior unless you explicitly specify otherwise by disabling rollback.

upvoted 3 times

🗨️ 👤 **Aamee** 7 months, 4 weeks ago

If you've an understanding of the concept of 'Circuit Breaker' for ECS based stacks, you'll come to know what exactly happens during the CFN failure. By default, the stack itself does not get deleted. It's the resources which get deleted IF and only IF you don't specify the 'Retain' configuration property in ur CFN script. Upon CFN failure, it normally shows 'ROLLBACK_COMPLETED' or 'ROLLBACK_IN_PROGRESS' that clearly shows that the stack of the Cloudformation script through which you're trying to create is not deleted.

upvoted 1 times

🗨️ 👤 **King_AWS_202** 1 year, 10 months ago

Selected Answer: B

CloudFormation will not delete the stack by default. If you want CloudFormation to delete the stack if it rolls back, you need to specify the DeletionPolicy property in the template. The DeletionPolicy property can have three values:

Delete - CloudFormation will delete the stack and all of its resources.

Retain - CloudFormation will not delete the stack or any of its resources.

Snapshot - CloudFormation will create a snapshot of the DB instance and then delete the stack.

In this case, the default behavior is to roll back the stack and not delete the stack.

upvoted 3 times

🗨️ 👤 **Andrew_A** 2 years ago

Selected Answer: B

CloudFormation will delete the resources it created during the stack creation attempt but it will not remove the stack record itself. This allows you to review the stack events and troubleshoot the reason for failure. I hope this clarifies the matter.

upvoted 4 times

🗨️ 👤 **Gomer** 2 years, 2 months ago

Second quote blurs the distinction between A and B, but I think answer is still B.

I actually creating a stack it with a template with an intentional error (acloud.guru demo/lab file).

The stack roll back out as expected and the stack did not delete itself, and I still had the stack listed under the "stacks" menu.

Even though it had known error and never had a "last known stable state", I saw it listed in stacks menu, and I wasn't able to create another stack using the "TestStack stack (e.g. Error "Stack name already exists")

Don't mean to ramble on here, but sometimes the AWS exam questions and answers don't exactly match the reality in the AWS console.

Facts:

(*) "Roll back all stack resources" (default setting from CloudFormation create stack menu)

"Specify the rollback-stack operation to roll back a stack to its last stable state."

"Note: The rollback-stack operation will delete a stack if it doesn't contain a last known stable state."

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stack-failure-options.html>

upvoted 3 times

🗨️ 👤 **SomboonCH** 2 years, 3 months ago

Selected Answer: A

I think A

The rollback-stack operation will delete a stack if it doesn't contain a last known stable state.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stack-failure-options.html>

upvoted 4 times

🗨️ 👤 **awsguru1998** 2 years, 4 months ago

B CloudFormation will roll back the stack in the event of a failed resource. It will not delete the stack automatically, and will instead leave it in a failed state. The user will then have the option to roll back the stack manually or continue with the stack in its failed state.

upvoted 2 times

🗨️ 👤 **anderri** 2 years, 4 months ago

A

"If the stack updates fails, CloudFormation rolls back changes to restore the stack to the last known working state."

Restore the stack means deleting the new objects created during the stack.

upvoted 1 times

🗨️ 👤 **defmania00** 2 years, 4 months ago

Yes, but it won't delete the stack, just de objects it created.

upvoted 1 times

🗨️ 👤 **Deeezz** 2 years, 5 months ago

The answer is A. Reference<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-what-is-howdoesitwork.html>

upvoted 1 times

🗨️ 👤 **Deeezz** 2 years, 5 months ago

"After all the resources have been created, CloudFormation reports that your stack has been created. You can then start using the resources in your stack. If stack creation fails, CloudFormation rolls back your changes by deleting the resources that it created." - AWS

upvoted 1 times

🗨️ 👤 **joanneli77** 2 years, 4 months ago

....annnnd the stack remains. Even with "Create Failed" status, you can highlight then delete the stack... this means the stack is not deleted. It's rolled back. If the stack is created the first time, the rollback is to no objects, but the stack remains, effectively empty. The answer is B.

upvoted 1 times

🗨️ 👤 **skywalker** 2 years, 5 months ago

Selected Answer: A

I think is A... it will attempt to delete the stack and rollback...

upvoted 2 times

A SysOps administrator needs to automate the invocation of an AWS Lambda function. The Lambda function must run at the end of each day to generate a report on data that is stored in an Amazon S3 bucket.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that has an event pattern for Amazon S3 and the Lambda function as a target.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that has a schedule and the Lambda function as a target.
- C. Create an S3 event notification to invoke the Lambda function whenever objects change in the S3 bucket.
- D. Deploy an Amazon EC2 instance with a cron job to invoke the Lambda function.

Suggested Answer: B

Community vote distribution

B (100%)

🗲️ 👤 **Liongeek** Highly Voted 👍 2 years, 1 month ago

Ans B

S3 Event notification has nothing to do here since the lambda will be triggered based on a daily schedule not every time an object is uploaded
upvoted 10 times

🗲️ 👤 **seetpt** Most Recent ⌚ 8 months, 4 weeks ago

Selected Answer: B

b is correct

upvoted 1 times

🗲️ 👤 **Rabbit117** 9 months, 4 weeks ago

B:

<https://aws.amazon.com/blogs/compute/introducing-amazon-eventbridge-scheduler/>

upvoted 2 times

🗲️ 👤 **michaldavid** 2 years ago

Selected Answer: B

bbbbbbbbb

upvoted 4 times

A company is releasing a new static website hosted on Amazon S3. The static website hosting feature was enabled on the bucket and content was uploaded; however, upon navigating to the site, the following error message is received:

403 Forbidden - Access Denied

What change should be made to fix this error?

- A. Add a bucket policy that grants everyone read access to the bucket.
- B. Add a bucket policy that grants everyone read access to the bucket objects.
- C. Remove the default bucket policy that denies read access to the bucket.
- D. Configure cross-origin resource sharing (CORS) on the bucket.

Suggested Answer: B

Community vote distribution

B (85%)

A (15%)

 **Domdom120** Highly Voted 2 years, 4 months ago

Selected Answer: B

B.

Per AWS: <https://repost.aws/knowledge-center/s3-static-website-endpoint-error>

"If an Access Denied error is returned by the web browser or cURL command, then the object isn't publicly accessible. To allow public read access to your S3 object, create a bucket policy that allows public read access for all objects in the bucket."

I'm not sure why everyone here is trying to provide answers without actual references. How you think in your head is irrelevant if it's not correct with AWS documentation.


upvoted 14 times

 **Liongeek** Highly Voted 2 years, 7 months ago

Ans B

Ref: <https://aws.amazon.com/es/premiumsupport/knowledge-center/s3-static-website-endpoint-error/>

upvoted 7 times

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

Answer: B

upvoted 1 times

 **csG13** 2 years, 4 months ago

Selected Answer: B

The answer is B, here is a sample bucket policy

```
{
  "Version": "2012-10-12",
  "Statement": {
    "Sid": "PublicReadGetObject",
    "Effect": "Allow",
    "Action": [
      "s3:GetObject"
    ],
    "Principal": "*",
    "Resource": [
      "arn:aws:s3:::example-s3-website.com/*"
    ]
  }
}
```

upvoted 6 times

🗨️ 👤 **Gil80** 2 years, 4 months ago

Selected Answer: B

It can't be A because Amazon doesn't have "Block Public Access Policy" feature, it has "Block Public Access Settings" - The key difference is Policy vs. Settings.

Therefore A is confusing to many.

Once you clear the "Block Public Access" SETTINGS you then ADD A BUCKET POLICY that makes your bucket content publicly available.

upvoted 2 times

🗨️ 👤 **Gil80** 2 years, 4 months ago

Selected Answer: B

This will allow anyone to access the objects in the bucket and view the static website.

Option A, "Add a bucket policy that grants everyone read access to the bucket," grants read access to the bucket itself, but not necessarily to its objects.

upvoted 1 times

🗨️ 👤 **joanneli77** 2 years, 4 months ago

Selected Answer: B

Reading objects, not just reading from the bucket, is the key to static website hosting. Nobody browses to the bucket, but instead the files themselves.

upvoted 3 times

🗨️ 👤 **skiwili** 2 years, 6 months ago

Selected Answer: B

Answer is B.

upvoted 1 times

🗨️ 👤 **MrMLB** 2 years, 6 months ago

Selected Answer: A

A. Add a bucket policy that grants everyone read access to the bucket.

To fix the error message "403 Forbidden - Access Denied" when trying to access a static website hosted on Amazon S3, you should add a bucket policy that grants everyone read access to the bucket. By default, Amazon S3 bucket policies deny all public access to the bucket and its contents. To allow users to access the static website hosted on the bucket, you need to add a bucket policy that grants read access to the bucket. You can do this by specifying a bucket policy that includes a "Principal" element with the value "*" (wildcard), which grants access to everyone, and an "Action" element with the value "s3:GetObject", which allows users to retrieve objects from the bucket.

Option B, adding a bucket policy that grants everyone read access to the bucket objects, would not be a valid solution, as this would not grant read access to the bucket itself.

upvoted 4 times

🗨️ 👤 **Aamee** 7 months, 4 weeks ago

COMPLETE WRONG IF YOU'RE GOING WITH OPTION A!....You're adding a bucket policy which grants Read access only to the bucket level?? Do you think everyone can read the bucket objects 'Magically' ??? Pls. read the question again and then you'll realize that the correct option is only B!

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: B

bbbbbb

upvoted 1 times

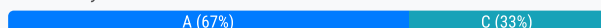
A company uses AWS Organizations. A SysOps administrator wants to use AWS Compute Optimizer and AWS tag policies in the management account to govern all member accounts in the billing family. The SysOps administrator navigates to the AWS Organizations console but cannot activate tag policies through the management account.

What could be the reason for this issue?

- A. All features have not been enabled in the organization.
- B. Consolidated billing has not been enabled.
- C. The member accounts do not have tags enabled for cost allocation.
- D. The member accounts have not manually enabled trusted access for Compute Optimizer.

Suggested Answer: A

Community vote distribution



🗨️ 👤 **Liongeek** Highly Voted 1 year, 7 months ago

Ans: A

Ref: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_tag-policies-prereqs.html

upvoted 10 times

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

ag policies and AWS Compute Optimizer are such features that need to be enabled explicitly within the organization.

upvoted 2 times

🗨️ 👤 **Hatem08** Most Recent 6 months, 3 weeks ago

Selected Answer: A

aaaaaa

upvoted 2 times

🗨️ 👤 **Hatem08** 7 months ago

Selected Answer: C

Ans: CCCC

upvoted 1 times

🗨️ 👤 **Hatem08** 7 months ago

sorry A is the correct

upvoted 3 times

🗨️ 👤 **King_AWS_202** 10 months ago

Selected Answer: C

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_org_support-all-features.html

upvoted 1 times

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

Selected Answer: A

aaaaaaa

upvoted 2 times

A company is storing media content in an Amazon S3 bucket and uses Amazon CloudFront to distribute the content to its users. Due to licensing terms, the company is not authorized to distribute the content in some countries. A SysOps administrator must restrict access to certain countries.

What is the MOST operationally efficient solution that meets these requirements?

- A. Configure the S3 bucket policy to deny the GetObject operation based on the S3:LocationConstraint condition.
- B. Create a secondary origin access identity (OAI). Configure the S3 bucket policy to prevent access from unauthorized countries.
- C. Enable the geo restriction feature in the CloudFront distribution to prevent access from unauthorized countries.
- D. Update the application to generate signed CloudFront URLs only for IP addresses in authorized counties.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Christina666** Highly Voted 1 year, 5 months ago

Selected Answer: C

You can use geographic restrictions, sometimes known as geo blocking, to prevent users in specific geographic locations from accessing content that you're distributing through a CloudFront distribution. To use geographic restrictions, you have two options:

Use the CloudFront geographic restrictions feature. Use this option to restrict access to all of the files that are associated with a distribution and to restrict access at the country level.

Use a third-party geolocation service. Use this option to restrict access to a subset of the files that are associated with a distribution or to restrict access at a finer granularity than the country level.

upvoted 6 times

🗳️ 👤 **Xelnak** Highly Voted 2 years, 1 month ago

Selected Answer: C

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

upvoted 5 times

🗳️ 👤 **Rabbit117** Most Recent 9 months, 4 weeks ago

Selected Answer: C

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

upvoted 2 times

🗳️ 👤 **pacheco00** 1 year, 3 months ago

You can use geographic restrictions, sometimes known as geo blocking, to prevent users in specific geographic locations from accessing content that you're distributing through a CloudFront distribution. To use geographic restrictions, you have two options:

Use the CloudFront geographic restrictions feature. Use this option to restrict access to all of the files that are associated with a distribution and to restrict access at the country level.

Use a third-party geolocation service. Use this option to restrict access to a subset of the files that are associated with a distribution or to restrict access at a finer granularity than the country level.

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: C

ccccc

upvoted 3 times

A SysOps administrator created an Amazon VPC with an IPv6 CIDR block, which requires access to the internet. However, access from the internet towards the VPC is prohibited. After adding and configuring the required components to the VPC, the administrator is unable to connect to any of the domains that reside on the internet.

What additional route destination rule should the administrator add to the route tables?

- A. Route `::/0` traffic to a NAT gateway
- B. Route `::/0` traffic to an internet gateway
- C. Route `0.0.0.0/0` traffic to an egress-only internet gateway
- D. Route `::/0` traffic to an egress-only internet gateway

Suggested Answer: D

Community vote distribution

D (82%)

B (18%)

 **tyfta6** Highly Voted 2 years, 6 months ago


Selected Answer: D

Vote for D

IPV4 = NAT Instance/Gateway | 0.0.0.0

IPV6 = Egress-Only Internet Gateway | `::/0`

upvoted 9 times

 **Liongeek** Highly Voted 2 years, 7 months ago

Ans: D

Ref: <https://docs.aws.amazon.com/vpc/latest/userguide/egress-only-internet-gateway.html>

upvoted 6 times

 **r2c3po** Most Recent 1 year, 6 months ago

Selected Answer: B

B. Route `::/0` traffic to an internet gateway

To enable IPv6 traffic from an Amazon VPC to the internet, you need to add a default route (`::/0`) pointing to an internet gateway in the VPC's route table. This allows outbound traffic to reach the internet.


This route configuration allows all IPv6 traffic (`::/0`) to be directed to the internet gateway, enabling communication with the internet.

Option A is incorrect because using a NAT gateway is typically for IPv4 traffic and is not applicable for IPv6.

Options C and D are not relevant for enabling outbound internet access in an IPv6-enabled VPC.

Egress-only internet gateways are used for allowing outbound traffic initiated by resources in the VPC to reach the internet, but they are not used for incoming traffic from the internet.

upvoted 1 times

 **Aamee** 9 months, 1 week ago

It's def. not 'B'. Read the question one more time. The VPC traffic only needs to communicate with the IGW and not vice versa. Therefore, D is correct considering the scenario discussed in this question IMO.

upvoted 1 times

 **Christina666** 1 year, 11 months ago

Selected Answer: D

IPv6 addresses are globally unique, and are therefore public by default. If you want your instance to be able to access the internet, but you want to prevent resources on the internet from initiating communication with your instance, you can use an egress-only internet gateway. To do this, create an egress-only internet gateway in your VPC, and then add a route to your route table that points all IPv6 traffic (`::/0`) or a specific range of IPv6 address to the egress-only internet gateway. IPv6 traffic in the subnet that's associated with the route table is routed to the egress-only internet gateway.

upvoted 4 times

🗨️ 👤 **Boul** 1 year, 11 months ago

It cannot be B, since access from the internet must be prohibited

upvoted 2 times

🗨️ 👤 **Cagarrieres** 2 years, 2 months ago

Dddddddddd

upvoted 1 times

🗨️ 👤 **skiwili** 2 years, 6 months ago

Selected Answer: D

Dddddddd

upvoted 4 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: D

ddddddd

upvoted 1 times

🗨️ 👤 **Xelnak** 2 years, 7 months ago

Selected Answer: B

B. Route ::/0 traffic to an internet gateway

NOT D because egress-only-internet-gateway is for accessing internet from private subnet

upvoted 2 times

🗨️ 👤 **beznika** 2 years, 6 months ago

It says that access from the internet is prohibited. So the only option correct here is D.

upvoted 2 times

🗨️ 👤 **jtz31** 2 years, 7 months ago

For me it's D.

We don't know if we try to access from private or public net, right?

upvoted 1 times

🗨️ 👤 **Aamee** 9 months, 1 week ago

The question doesn't state specifically about either public or private subnets here. It clearly states this "access from the internet towards the VPC is prohibited". That means the traffic from VPC need to communicate with IGW to the internet via an egress route option only and not vice versa. Therefore D is absolutely satisfying this requirement.

upvoted 1 times

🗨️ 👤 **fig** 1 year, 7 months ago

IPv6 is Public by nature... So it is D

upvoted 1 times

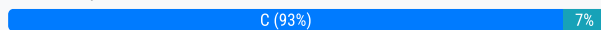
A company hosts several write-intensive applications. These applications use a MySQL database that runs on a single Amazon EC2 instance. The company asks a SysOps administrator to implement a highly available database solution that is ideal for multi-tenant workloads.

Which solution should the SysOps administrator implement to meet these requirements?

- A. Create a second EC2 instance for MySQL. Configure the second instance to be a read replica.
- B. Migrate the database to an Amazon Aurora DB cluster. Add an Aurora Replica.
- C. Migrate the database to an Amazon Aurora multi-master DB cluster.
- D. Migrate the database to an Amazon RDS for MySQL DB instance.

Suggested Answer: C

Community vote distribution



🗳️ 👤 **Xelnak** Highly Voted 2 years, 7 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-multi-master.html#aurora-multi-master-overview>

Multi-master clusters are best suited for segmented workloads, such as for multitenant applications.

upvoted 8 times

🗳️ 👤 **xdkonorek2** Highly Voted 1 year, 5 months ago

question is deprecated, there is no multi-master mode in aurora mysql anymore

upvoted 6 times

🗳️ 👤 **XXXXXINN** Most Recent 7 months, 3 weeks ago

C.

keyword "write-intensive" in the questions, that leads to multi-master DB...

upvoted 1 times

🗳️ 👤 **10cc6ba** 11 months, 3 weeks ago

Selected Answer: B

B

not C because

C. Migrate the database to an Amazon Aurora multi-master DB cluster.

While this provides high availability and is designed for write-intensive workloads, it's generally more complex and may be more costly. It's suitable for specific use cases where you need write scaling across multiple masters, which might not be necessary here.

upvoted 1 times

🗳️ 👤 **james2033** 1 year, 4 months ago

Selected Answer: C

Question's keyword 'write-intense' . Answer's keyword 'multi-master' , then choose Answer C.

upvoted 2 times

🗳️ 👤 **r2c3po** 1 year, 6 months ago

Selected Answer: C

C. Migrate the database to an Amazon Aurora multi-master DB cluster:

Amazon Aurora multi-master enables multiple instances to write to the database simultaneously. This can provide high availability and fault tolerance, making it suitable for write-intensive workloads.

It supports read and write operations across multiple instances, which is beneficial for multi-tenant workloads where different tenants may have different write patterns.



upvoted 2 times

🗳️ 👤 **alexiscloud** 1 year, 7 months ago

multi-tenant workload

Answer: C



upvoted 2 times

  **jipark** 1 year, 10 months ago

Selected Answer: C

multi-master = multi-tenant workload

upvoted 4 times

  **Liongeek** 2 years, 7 months ago

Ans: C

Ref: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-multi-master.html#aurora-multi-master-overview>

upvoted 4 times

A company has a memory-intensive application that runs on a fleet of Amazon EC2 instances behind an Elastic Load Balancer (ELB). The instances run in an Auto Scaling group. A SysOps administrator must ensure that the application can scale based on the number of users that connect to the application.

Which solution will meet these requirements?

- A. Create a scaling policy that will scale the application based on the ActiveConnectionCount Amazon CloudWatch metric that is generated from the ELB.
- B. Create a scaling policy that will scale the application based on the mem_used Amazon CloudWatch metric that is generated from the ELB.
- C. Create a scheduled scaling policy to increase the number of EC2 instances in the Auto Scaling group to support additional connections.
- D. Create and deploy a script on the ELB to expose the number of connected users as a custom Amazon CloudWatch metric. Create a scaling policy that uses the metric.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **r2c3po** 1 year ago

Selected Answer: A

Option A is the correct choice:

A. Create a scaling policy that will scale the application based on the ActiveConnectionCount CloudWatch metric:

The ActiveConnectionCount metric is a measure of the current number of open connections to the ELB. Scaling based on this metric allows the Auto Scaling group to adjust the number of instances based on the demand reflected in the number of active connections.

upvoted 1 times

🗳️ 👤 **alexiscloud** 1 year, 1 month ago

the key is to publish a CloudWatch metric that tracks active connections, and configure Auto Scaling scaling policies based on that metric to dynamically scale the EC2 fleet behind the ELB as the load changes.

Ans: A

upvoted 2 times

🗳️ 👤 **jipark** 1 year, 4 months ago

Selected Answer: A

must ensure that the application can scale

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: A

aaaaaaa

upvoted 2 times

🗳️ 👤 **marcelodba** 2 years, 1 month ago

Selected Answer: A

Ans. A

upvoted 3 times

A SysOps administrator creates a new VPC that includes a public subnet and a private subnet. The SysOps administrator successfully launches 11 Amazon EC2 instances in the private subnet. The SysOps administrator attempts to launch one more EC2 instance in the same subnet. However, the SysOps administrator receives an error message that states that not enough free IP addresses are available.

What must the SysOps administrator do to deploy more EC2 instances?

- A. Edit the private subnet to change the CIDR block to /27.
- B. Edit the private subnet to extend across a second Availability Zone.
- C. Assign additional Elastic IP addresses to the private subnet.
- D. Create a new private subnet to hold the required EC2 instances.

Suggested Answer: D

Community vote distribution

D (85%)

A (15%)

🗳️ 👤 **Christina666** Highly Voted 1 year, 5 months ago

Selected Answer: D

It's not possible to change or modify the IP address range of an existing virtual private cloud (VPC) or subnet. However, you can do one of the following:

Add an additional IPv4 CIDR block as a secondary CIDR to your VPC.

Create a new VPC with your preferred CIDR block and then migrate the resources from your old VPC to the new VPC (if applicable).

upvoted 12 times

🗳️ 👤 **Liongeek** Highly Voted 2 years, 1 month ago

Ans: D

Ref: <https://aws.amazon.com/es/premiumsupport/knowledge-center/vpc-ip-address-range/>

upvoted 7 times

🗳️ 👤 **Nomzie** Most Recent 10 months, 4 weeks ago

Selected Answer: D

It's not possible to change or modify the IP address range of an existing virtual private cloud (VPC) or subnet. However, you can do one of the following:

Add an additional IPv4 CIDR block as a secondary CIDR to your VPC.

Create a new VPC with your preferred CIDR block and then migrate the resources from your old VPC to the new VPC (if applicable).

Source: <https://repost.aws/es/knowledge-center/vpc-ip-address-range>

upvoted 2 times

🗳️ 👤 **r2c3po** 1 year ago

Selected Answer: A

A is the correct choice:

A. Edit the private subnet to change the CIDR block to /27:

By increasing the CIDR block size to /27 or a larger range, you allocate more IP addresses to the subnet, allowing for the deployment of additional EC2 instances.

Options B, C, and D are not appropriate solutions for addressing the issue:

B. Editing the private subnet to extend across a second Availability Zone does not directly resolve the lack of available IP addresses in the existing subnet.

C. Assigning additional Elastic IP addresses to the private subnet is not a solution for expanding the available private IP addresses within the subnet. Elastic IP addresses are public IP addresses associated with instances.

D. Creating a new private subnet is a valid option, but it involves additional steps such as moving resources, updating security groups, and reconfiguring routing. Editing the existing subnet's CIDR block is a more straightforward solution.

upvoted 1 times

🗨️ 👤 **Grodgar** 6 months, 1 week ago

Wrong, CIDR block can't be edited

upvoted 1 times

🗨️ 👤 **fig** 1 year, 1 month ago

Selected Answer: D

You cannot change the CIDR block once created. So not A. Answer is D

upvoted 4 times

🗨️ 👤 **johnnyjin** 1 year, 3 months ago

Selected Answer: A

Should be A, from /28(16) to /27(32) then the subnet will have more IP address available

upvoted 2 times

🗨️ 👤 **fig** 1 year, 1 month ago

You cannot change the CIDR block once created. So not A. Answer is D

upvoted 2 times

🗨️ 👤 **goatbernard** 2 years ago

Selected Answer: D

D is the answer

upvoted 3 times

A company needs to automatically monitor an AWS account for potential unauthorized AWS Management Console logins from multiple geographic locations.

Which solution will meet this requirement?

- A. Configure Amazon Cognito to detect any compromised IAM credentials.
- B. Set up Amazon Inspector. Scan and monitor resources for unauthorized logins.
- C. Set up AWS Config. Add the iam-policy-blacklisted-check managed rule to the account.
- D. Configure Amazon GuardDuty to monitor the UnauthorizedAccess:IAMUser/ConsoleLoginSuccess.B finding.

Suggested Answer: D

Community vote distribution

D (100%)

 **Christina666** Highly Voted 1 year, 5 months ago

Selected Answer: D

Guard duty IAM finding types:

UnauthorizedAccess:IAMUser/ConsoleLoginSuccess.B

UnauthorizedAccess:IAMUser/InstanceCredentialExfiltration.InsideAWS

UnauthorizedAccess:IAMUser/InstanceCredentialExfiltration.OutsideAWS

UnauthorizedAccess:IAMUser/MaliciousIPCaller

UnauthorizedAccess:IAMUser/MaliciousIPCaller.Custom

UnauthorizedAccess:IAMUser/TorIPCaller

upvoted 7 times

 **r2c3po** Most Recent 1 year ago

Selected Answer: D

GuardDuty provides findings related to unauthorized access and console logins. The "UnauthorizedAccess:IAMUser/ConsoleLoginSuccess.B" finding specifically indicates successful console logins from new or unusual locations.

By configuring GuardDuty, you can receive alerts and notifications for such findings, allowing you to investigate and take appropriate action.

Options A, B, and C are not directly designed for monitoring unauthorized console logins from multiple geographic locations:

A. Amazon Cognito is an identity management service but is not specifically designed for monitoring AWS Management Console logins.

B. Amazon Inspector is a security assessment service that identifies security vulnerabilities in your EC2 instances. While it can help with overall security posture, it is not focused on monitoring console logins.

C. AWS Config can provide information about configuration changes, but the "iam-policy-blacklisted-check" managed rule is related to IAM policies and is not specifically designed for monitoring console logins from multiple locations.

upvoted 2 times

 **alexiscloud** 1 year, 1 month ago

The main things GuardDuty monitors for are:

Unauthorized API calls

IAM policy anomalies

Resource inconsistencies

Resource usage anomalies

Ans:D

upvoted 2 times

 **Christina666** 1 year, 5 months ago

Selected Answer: D

Default severity: Medium

Data source: CloudTrail management events

This finding informs you that multiple successful console logins for the same IAM user were observed around the same time in various geographical locations. Such anomalous and risky access location patterns indicate potential unauthorized access to your AWS resources.

upvoted 3 times

🗨️ 👤 **tts1234** 2 years ago

Selected Answer: D

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_finding-types-iam.html#unauthorizedaccess-iam-consoleloginsuccessb

upvoted 3 times

🗨️ 👤 **michaldavid** 2 years ago

Selected Answer: D

ddddddd

upvoted 3 times

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

Ans: D

upvoted 2 times

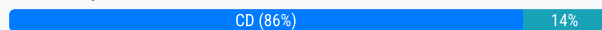
A company has an Amazon RDS DB instance. The company wants to implement a caching service while maintaining high availability.

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

- A. Add Auto Discovery to the data store.
- B. Create an Amazon ElastiCache for Memcached data store.
- C. Create an Amazon ElastiCache for Redis data store.
- D. Enable Multi-AZ for the data store.
- E. Enable Multi-threading for the data store.

Suggested Answer: CD

Community vote distribution



Liongeek Highly Voted 2 years, 1 month ago

Ans: C&D

C because only Redis gives cache and HA (Memcached doesn't have HA)

Ref: <https://aws.amazon.com/es/elasticache/redis-vs-memcached/>

D: For HA

upvoted 12 times

henro4niger Most Recent 10 months, 4 weeks ago

Answer is definitely CD. Memcached does not support multi-az deployments which is required for HA as stated in the question. Performance was not necessarily in the question, hence, multi-threading is not required.

CD for sure

upvoted 3 times

r2c3po 1 year ago

Selected Answer: BD

B. Create an Amazon ElastiCache for Memcached data store.

D. Enable Multi-AZ for the data store.

To implement a caching service with high availability, you can use Amazon ElastiCache, which is a fully managed caching service compatible with both Memcached and Redis. For high availability, you should enable Multi-AZ (Availability Zone) support.

#B. Create an Amazon ElastiCache for Memcached data store:

Memcached is a widely used caching system, and ElastiCache for Memcached is suitable for simple caching scenarios. Choose this option if Memcached meets your caching requirements.

#D. Enable Multi-AZ for the data store:

Enabling Multi-AZ for Amazon RDS or Amazon ElastiCache ensures that the service is deployed across multiple Availability Zones, providing redundancy and high availability. In case of a failure in one Availability Zone, the service continues to operate in another.

upvoted 1 times

alexiscloud 1 year, 1 month ago

The two options that will best meet the requirements are:

B. Create an Amazon ElastiCache for Memcached data store.

D. Enable Multi-AZ for the data store.

Summary:

Creating an Amazon ElastiCache Memcached data store will provide a caching layer in front of the RDS DB instance. This will improve performance

by caching frequently accessed data.

Enabling Multi-AZ for the RDS DB instance will maintain high availability. The Multi-AZ configuration will replicate the DB across Availability Zones, so if one AZ fails, the DB will still be accessible from the other AZ.

Answer:B,D

upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: CD

C and D

upvoted 3 times

  **marcelodba** 2 years ago

Selected Answer: CD

multi-thread is related to performance.

<https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html#AutoFailover.Enable>

upvoted 3 times

A company monitors its account activity using AWS CloudTrail, and is concerned that some log files are being tampered with after the logs have been delivered to the account's Amazon S3 bucket.

Moving forward, how can the SysOps administrator confirm that the log files have not been modified after being delivered to the S3 bucket?

- A. Stream the CloudTrail logs to Amazon CloudWatch Logs to store logs at a secondary location.
- B. Enable log file integrity validation and use digest files to verify the hash value of the log file.
- C. Replicate the S3 log bucket across regions, and encrypt log files with S3 managed keys.
- D. Enable S3 server access logging to track requests made to the log bucket for security audits.

Suggested Answer: B

Community vote distribution



B (100%)

  **JamesF92**  1 year, 4 months ago

Selected Answer: B

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

upvoted 5 times

  **jipark** 1 year, 4 months ago

"integrity validation" is keyword

upvoted 1 times

  **Liongeek**  2 years, 1 month ago

Ans: B

upvoted 5 times

  **r2c3po**  1 year ago

Selected Answer: B

Option B is the correct choice:

To confirm that CloudTrail log files have not been modified after being delivered to an S3 bucket, you can enable log file integrity validation. When log file integrity validation is enabled, AWS CloudTrail generates digest files that contain the hash values of the delivered log files. These hash values are then used to verify the integrity of the log files.

B. Enable log file integrity validation and use digest files to verify the hash value of the log file:

Log file integrity validation helps ensure that log files stored in the S3 bucket have not been tampered with. It adds an additional layer of security by providing a way to verify the integrity of the log files using hash values stored in digest files.

upvoted 2 times

  **Christina666** 1 year, 5 months ago

Selected Answer: B

B.....log file integrity validation

upvoted 2 times

  **michaldavid** 2 years ago

Selected Answer: B

bbbbbbb

upvoted 3 times

A SysOps administrator is reviewing AWS Trusted Advisor warnings and encounters a warning for an S3 bucket policy that has open access permissions. While discussing the issue with the bucket owner, the administrator realizes the S3 bucket is an origin for an Amazon CloudFront web distribution.

Which action should the administrator take to ensure that users access objects in Amazon S3 by using only CloudFront URLs?

- A. Encrypt the S3 bucket content with Server-Side Encryption with Amazon S3-Managed Keys (SSE-S3).
- B. Create an origin access identity and grant it permissions to read objects in the S3 bucket.
- C. Assign an IAM user to the CloudFront distribution and grant the user permissions in the S3 bucket policy.
- D. Assign an IAM role to the CloudFront distribution and grant the role permissions in the S3 bucket policy.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **r2c3po** 1 year ago

Selected Answer: B

Option B is the correct choice:

#B. Create an origin access identity and grant it permissions to read objects in the S3 bucket.

When using Amazon CloudFront as a content delivery network with an S3 bucket as the origin, it's a best practice to restrict direct access to the S3 bucket and require users to access objects only through CloudFront URLs.

This can be achieved by creating an Origin Access Identity (OAI) and granting it permission to read objects in the S3 bucket.

upvoted 3 times

🗳️ 👤 **Saibal9** 1 year ago

It is actually Origin Access Control now.

upvoted 2 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: B

OAI-----only Cloudfront get objects from S3

upvoted 4 times

🗳️ 👤 **jipark** 1 year, 4 months ago

got it, OAI is used for CloudFront

upvoted 2 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: B

bbbbbb

upvoted 3 times

🗳️ 👤 **Liongeek** 2 years, 1 month ago

Ans: B

I just had a class on that in CloudGuru :p

upvoted 3 times

A SysOps administrator is reviewing AWS Trusted Advisor recommendations. The SysOps administrator notices that all the application servers for a finance application are listed in the Low Utilization Amazon EC2 Instances check. The application runs on three instances across three Availability Zones. The SysOps administrator must reduce the cost of running the application without affecting the application's availability or design.

Which solution will meet these requirements?

- A. Reduce the number of application servers.
- B. Apply rightsizing recommendations from AWS Cost Explorer to reduce the instance size.
- C. Provision an Application Load Balancer in front of the instances.
- D. Scale up the instance size of the application servers.

Suggested Answer: B

Community vote distribution

B (100%)

  **Gomer** Highly Voted 8 months, 2 weeks ago

Selected Answer: B

For what it's worth, since the app servers are installed in three different AZ's, you could drop an app server in one AZ and meet the HA requirement. However, this does make some presumptions on the utilization. However, I do believe B is probably the right answer (from researching AWS documentation)

"The rightsizing recommendations feature in Cost Explorer helps you identify cost-saving opportunities by downsizing or terminating instances in Amazon Elastic Compute Cloud (Amazon EC2). Rightsizing recommendations analyze your Amazon EC2 resources and usage to show opportunities for how you can lower your spending. You can see all of your underutilized Amazon EC2 instances across member accounts in a single view to immediately identify how much you can save. After you identify your recommendations, you can take action on the Amazon EC2 console."

<https://docs.aws.amazon.com/cost-management/latest/userguide/ce-rightsizing.html>

upvoted 5 times

  **beznika** Most Recent 1 year ago

B for sure. You can't reduce the number of instances.

upvoted 2 times

  **michaldavid** 1 year ago

Selected Answer: B

bbbbbbbb

upvoted 1 times

  **Fatoch** 1 year ago

is this B? Anyone have idea?

upvoted 2 times

  **gsotiriou** 1 year ago

Well... For me it's is certainly not:

C -> LoadBalancing will not reduce your cost in the end

D -> Increasing the size of your servers is pretty much obvious why is not a good idea

That leaves A or B.

I believe it is not A because the use case tells us to NOT mess with the application's "high availability and design". That means that the application should still be running on a 3-server configuration across 3 availability zones as that ensures HA and might have special design considerations in its programming.

So that leaves us with B and trusting the AWS Cost Explorer and its right sizing.

upvoted 1 times

A company hosts its website in the us-east-1 Region. The company is preparing to deploy its website into the eu-central-1 Region. Website visitors who are located in Europe should access the website that is hosted in eu-central-1. All other visitors access the website that is hosted in us-east-1. The company uses Amazon Route 53 to manage the website's DNS records.

Which routing policy should a SysOps administrator apply to the Route 53 record set to meet these requirements?

- A. Geolocation routing policy
- B. Geoproximity routing policy
- C. Latency routing policy
- D. Multivalue answer routing policy

Suggested Answer: A

Community vote distribution

A (100%)

  **edu_anadat** Highly Voted 2 years ago



Selected Answer: A

A

Geolocation routing policy --> it's used to user location

Geoproximity routing policy --> it's used to resource location

upvoted 14 times

  **jipark** 1 year, 4 months ago

Geoproximity decide route based on capacity / resource of location. (not distance)

upvoted 2 times

  **Rabbit117** Most Recent 9 months, 4 weeks ago

Selected Answer: A

Geolocation routing lets you choose the resources that serve your traffic based on the geographic location of your users, meaning the location that DNS queries originate from. For example, you might want all queries from Europe to be routed to an Elastic Load Balancing load balancer in the Frankfurt Region.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-geo.html>

upvoted 2 times

  **r2c3po** 1 year ago

Selected Answer: A

Option A is the correct choice: A. Geolocation routing policy:

With Geolocation routing, you can specify different DNS records based on the geographic location of the user. In this case, you would create a record set for the eu-central-1 Region and another for the us-east-1 Region.

Option B, Geoproximity routing policy, is designed for routing traffic based on the geographic location of the resources, not the users.


upvoted 3 times

  **grka25** 2 years ago

I think this should be B

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-geoproximity.html>

upvoted 1 times

  **xSohox** 1 year, 4 months ago

It should be A.

You create 2 records with Geolocation routing policy:

1 is for Europe location

2 is Default

upvoted 3 times

  **michaldavid** 2 years ago

Selected Answer: A

aaaaaaa

upvoted 2 times

  **Fatoch** 2 years ago

It's A

upvoted 1 times

An organization with a large IT department has decided to migrate to AWS. With different job functions in the IT department, it is not desirable to give all users access to all AWS resources. Currently the organization handles access via LDAP group membership.

What is the BEST method to allow access using current LDAP credentials?

- A. Create an AWS Directory Service Simple AD. Replicate the on-premises LDAP directory to Simple AD.
- B. Create a Lambda function to read LDAP groups and automate the creation of IAM users.
- C. Use AWS CloudFormation to create IAM roles. Deploy Direct Connect to allow access to the on-premises LDAP server.
- D. Federate the LDAP directory with IAM using SAML. Create different IAM roles to correspond to different LDAP groups to limit permissions.

Suggested Answer: D

Community vote distribution

D (100%)

  **tyfta6**  2 years ago

Selected Answer: D

Vote for D

upvoted 5 times

  **johnson_chao**  9 months ago

Selected Answer: D

answer is D

ref : https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_enable-console-saml.html

upvoted 2 times

 r2c3po 1 year ago

Selected Answer: D

Option D is the correct choice:

D. Federate the LDAP directory with IAM using SAML:

By federating with IAM using SAML, you enable single sign-on (SSO) and allow users to authenticate using their existing LDAP credentials. This eliminates the need to manage separate IAM user accounts and passwords for AWS access.

Different IAM roles can be created to correspond to different LDAP groups, and users can assume these roles to gain access to AWS resources. This helps in limiting permissions based on the LDAP group membership.

upvoted 2 times

 alexiscloud 1 year, 1 month ago

Uses existing LDAP credentials

Allows role mapping by LDAP group

Fulfills access requirements by creating granular IAM roles mapped to LDAP groups

Ans:D

upvoted 2 times

 grka25 2 years ago

Answer is D

upvoted 4 times

A SysOps administrator has created an Amazon EC2 instance using an AWS CloudFormation template in the us-east-1 Region. The administrator finds that this template has failed to create an EC2 instance in the us-west-2 Region.

What is one cause for this failure?

- A. Resource tags defined in the CloudFormation template are specific to the us-east-1 Region.
- B. The Amazon Machine Image (AMI) ID referenced in the CloudFormation template could not be found in the us-west-2 Region.
- C. The cfn-init script did not run during resource provisioning in the us-west-2 Region.
- D. The IAM user was not created in the specified Region.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **r2c3po** 1 year ago

Selected Answer: B

B is the correct choice:

B. The Amazon Machine Image (AMI) ID referenced in the CloudFormation template could not be found in the us-west-2 Region:

AMIs are region-specific, and the AMI ID referenced in the template must be valid in the target region (us-west-2, in this case). If the specified AMI is not available in the us-west-2 Region, the CloudFormation stack creation will fail.

upvoted 3 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: B

failed to create EC2-> AMI may not exist as it's region specific

upvoted 3 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: B

bbbbbbbbb

upvoted 1 times

🗳️ 👤 **tyfta6** 2 years ago

Selected Answer: B

Vote for B

AMI is region specific. So it might not be available in another region.

upvoted 2 times

🗳️ 👤 **grka25** 2 years ago

Answer is B

upvoted 2 times

A global gaming company is preparing to launch a new game on AWS. The game runs in multiple AWS Regions on a fleet of Amazon EC2 instances. The instances are in an Auto Scaling group behind an Application Load Balancer (ALB) in each Region. The company plans to use Amazon Route 53 for DNS services. The DNS configuration must direct users to the Region that is closest to them and must provide automated failover.

Which combination of steps should a SysOps administrator take to configure Route 53 to meet these requirements? (Choose two.)

- A. Create Amazon CloudWatch alarms that monitor the health of the ALB in each Region. Configure Route 53 DNS failover by using a health check that monitors the alarms.
- B. Create Amazon CloudWatch alarms that monitor the health of the EC2 instances in each Region. Configure Route 53 DNS failover by using a health check that monitors the alarms.
- C. Configure Route 53 DNS failover by using a health check that monitors the private IP address of an EC2 instance in each Region.
- D. Configure Route 53 geoproximity routing. Specify the Regions that are used for the infrastructure.
- E. Configure Route 53 simple routing. Specify the continent, country, and state or province that are used for the infrastructure.

Suggested Answer: AD

Community vote distribution

AD (85%)

CD (15%)

  **Gil80** Highly Voted 1 year, 10 months ago

Selected Answer: AD

Option B is not correct because monitoring the health of the EC2 instances is not sufficient to provide failover as the EC2 instances are in an Auto Scaling group and instances can be added or removed dynamically.

Option C is not correct because monitoring the private IP address of an EC2 instance is not sufficient to determine the health of the infrastructure, as the instance may still be running but the application or service on the instance may be unhealthy.

Option E is not correct because simple routing does not take into account geographic proximity, which is a requirement in this scenario.

upvoted 10 times

  **jipark** Highly Voted 1 year, 4 months ago

Selected Answer: AD

two clues are :

1. "EC2 behind ALB" -> monitor ALB
2. route "users" -> geoproximity

upvoted 6 times

  **r2c3po** Most Recent 1 year ago

Selected Answer: AD

A. Create Amazon CloudWatch alarms that monitor the health of the ALB in each Region. Configure Route 53 DNS failover by using a health check that monitors the alarms.

D. Configure Route 53 geoproximity routing. Specify the Regions that are used for the infrastructure.

upvoted 2 times

  **Phinx** 1 year, 10 months ago

Selected Answer: CD

C and D. No other options. CloudWatch is irrelevant in this scenario.

upvoted 4 times

  **Kipalom** 1 year ago

You can't use Route 53 DNS Failover for private IPs

upvoted 2 times

  **Bar_t** 10 months, 3 weeks ago

"Route 53 health checkers are outside the VPC. To check the health of an endpoint within a VPC by IP address, you must assign a public IP address to the instance in the VPC."

upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: AD

A and D

upvoted 6 times

A SysOps administrator is investigating a company's web application for performance problems. The application runs on Amazon EC2 instances that are in an Auto Scaling group. The application receives large traffic increases at random times throughout the day. During periods of rapid traffic increases, the Auto Scaling group is not adding capacity fast enough. As a result, users are experiencing poor performance.

The company wants to minimize costs without adversely affecting the user experience when web traffic surges quickly. The company needs a solution that adds more capacity to the Auto Scaling group for larger traffic increases than for smaller traffic increases.

How should the SysOps administrator configure the Auto Scaling group to meet these requirements?

- A. Create a simple scaling policy with settings to make larger adjustments in capacity when the system is under heavy load.
- B. Create a step scaling policy with settings to make larger adjustments in capacity when the system is under heavy load.
- C. Create a target tracking scaling policy with settings to make larger adjustments in capacity when the system is under heavy load.
- D. Use Amazon EC2 Auto Scaling lifecycle hooks. Adjust the Auto Scaling group's maximum number of instances after every scaling event.

Suggested Answer: B

Community vote distribution

B (69%)

C (28%)

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

Selected Answer: B

I have to vote "B" because if you continue to read the whole quote, AWS recommends step scaling over target tracking in this instance, e.g.:

"We strongly recommend that you use a target tracking scaling policy to scale on a metric like average CPU utilization or the RequestCountPerTarget metric from the Application Load Balancer. Metrics that decrease when capacity increases and increase when capacity decreases can be used to proportionally scale out or in the number of instances using target tracking. This helps ensure that Amazon EC2 Auto Scaling follows the demand curve for your applications closely."

...

"You still have the option to use step scaling as an additional policy for a more advanced configuration. For example, you can configure a more aggressive response when demand reaches a certain level."

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html>

upvoted 14 times

 **dangji** Highly Voted 1 year, 11 months ago

Selected Answer: C

C AWS recommend use "target tracking scaling".

"We strongly recommend that you use a target tracking scaling policy to scale on a metric like average CPU utilization or the RequestCountPerTarget metric from the Application Load Balancer. "

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html>

upvoted 11 times

 **r2c3po** Most Recent 1 year ago

Selected Answer: B

B is the correct choice:

B. Create a step scaling policy with settings to make larger adjustments in capacity when the system is under heavy load:

Step scaling policies allow you to configure different scaling adjustments for different breach levels. In this case, you can set larger adjustments for larger increases in traffic. This helps the Auto Scaling group add more capacity quickly during periods of rapid traffic increases.

upvoted 2 times

 **Hatem08** 1 year ago

Selected Answer: B

B is the correct one!

upvoted 2 times

 **xile1021** 1 year, 2 months ago

Selected Answer: B

B, the question says "During periods of rapid traffic increases" meaning more than one increase.

"In contrast, with step scaling the policy can continue to respond to additional alarms, even while a scaling activity or health check replacement is in progress. Therefore, all alarms that are breached are evaluated by Amazon EC2 Auto Scaling as it receives the alarm messages."

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html#as-step-scaling-warmup>

upvoted 2 times

🗳️ 👤 **mayoness** 1 year, 2 months ago

"The main difference between the policy types is the step adjustments that you get with step scaling policies. When step adjustments are applied, and they increase or decrease the current capacity of your Auto Scaling group, the adjustments vary based on the size of the alarm breach." -- C

upvoted 1 times

🗳️ 👤 **Dafukubai** 1 year, 4 months ago

Selected Answer: B

B for me,

"You still have the option to use step scaling as an additional policy for a more advanced configuration. For example, you can configure a more aggressive response when demand reaches a certain level."

upvoted 3 times

🗳️ 👤 **delak** 1 year, 4 months ago

Selected Answer: C

C for me,

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

upvoted 2 times

🗳️ 👤 **jipark** 1 year, 4 months ago

Selected Answer: B

why not C: Target Scaling must wait for additional alarms (cause delay)

why B: Step Scaling handle multiple alarms.

upvoted 4 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: B

The correct answer is B. Create a step scaling policy with settings to make larger adjustments in capacity when the system is under heavy load.

Step scaling policies are suitable for scenarios where you want to add more capacity to the Auto Scaling group as the traffic increases. With step scaling policies, you can define specific thresholds for different levels of traffic and configure the scaling adjustments accordingly. This allows you to add more capacity during larger traffic increases while minimizing costs during smaller traffic increases.

Option A (Create a simple scaling policy) might not be the best choice as it does not provide the granularity needed to handle larger traffic increases differently from smaller ones.

upvoted 8 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Option C (Create a target tracking scaling policy) sets a target metric (e.g., CPU utilization or request count per instance) and automatically adjusts the capacity to maintain that target. While it can be useful for maintaining a specific performance level, it may not directly address the requirement of adding more capacity during large traffic increases.

Option D (Use Amazon EC2 Auto Scaling lifecycle hooks and adjust the maximum number of instances after every scaling event) is not the most efficient approach for handling traffic surges and may not be suitable for minimizing costs while maintaining performance.

Therefore, option B is the best choice for configuring the Auto Scaling group to handle rapid traffic increases effectively and minimize costs without adversely affecting the user experience.

upvoted 7 times

🗳️ 👤 **TQM_9MD** 1 year, 6 months ago

Selected Answer: C

C is answer

upvoted 1 times

🗳️ 👤 **braveheart22** 1 year, 8 months ago

BBB is the right answer and not C

In contrast, with step scaling the policy can continue to respond to additional alarms, even while a scaling activity or health check replacement is in

progress. Therefore, all alarms that are breached are evaluated by Amazon EC2 Auto Scaling as it receives the alarm messages.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html#SimpleScaling>

upvoted 3 times

🗲️ 👤 **noahsark** 1 year, 9 months ago

Selected Answer: B

When step adjustments are applied, and they increase or decrease the current capacity of your Auto Scaling group, the adjustments vary based on the size of the alarm breach.

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html>

upvoted 5 times

🗲️ 👤 **Vivec** 1 year, 9 months ago

Selected Answer: C

With a target tracking scaling policy, the SysOps administrator can define a target value for a specific metric that the Auto Scaling group will try to maintain. When the metric exceeds the target value, the policy will automatically add more instances to the Auto Scaling group. By configuring the scaling policy to make larger adjustments in capacity when the system is under heavy load, the Auto Scaling group will add more capacity for larger traffic increases. This will help to improve the user experience during periods of rapid traffic increases, without incurring unnecessary costs during normal traffic periods.

upvoted 4 times

🗲️ 👤 **joanneli77** 1 year, 11 months ago

Answer is "B" - you need to identify a specific metric to use "C" and there's no indication that is a "known quantity." Both answers are very close, but I have to go with "B" for this reason.

upvoted 2 times

🗲️ 👤 **Jaro3000** 1 year, 11 months ago

Selected Answer: B

Looks like B

upvoted 5 times

🗲️ 👤 **tts1234** 2 years ago

Selected Answer: B

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html>

upvoted 5 times

A company has a compliance requirement that no security groups can allow SSH ports to be open to all IP addresses. A SysOps administrator must implement a solution that will notify the company's SysOps team when a security group rule violates this requirement. The solution also must remediate the security group rule automatically.

Which solution will meet these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function when a security group changes. Configure the Lambda function to evaluate the security group for compliance, remove all inbound security group rules on all ports, and notify the SysOps team if the security group is noncompliant.
- B. Create an AWS CloudTrail metric filter for security group changes. Create an Amazon CloudWatch alarm to notify the SysOps team through an Amazon Simple Notification Service (Amazon SNS) topic when the metric is greater than 0. Subscribe an AWS Lambda function to the SNS topic to remediate the security group rule by removing the rule.
- C. Activate the AWS Config restricted-ssh managed rule. Add automatic remediation to the AWS Config rule by using the AWS Systems Manager Automation AWS-DisablePublicAccessForSecurityGroup runbook. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to notify the SysOps team when the rule is noncompliant.
- D. Create an AWS CloudTrail metric filter for security group changes. Create an Amazon CloudWatch alarm for when the metric is greater than 0. Add an AWS Systems Manager action to the CloudWatch alarm to suspend the security group by using the Systems Manager Automation AWS-DisablePublicAccessForSecurityGroup runbook when the alarm is in ALARM state. Add an Amazon Simple Notification Service (Amazon SNS) topic as a second target to notify the SysOps team.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **28b8844** 3 months, 1 week ago

Selected Answer: D

Option C is incomplete. It does not mention using AWS SNS for notification which is one of the key requirement.
upvoted 1 times

🗳️ 👤 **r2c3po** 1 year ago

Selected Answer: C

This solution combines AWS Config for rule evaluation, AWS Systems Manager Automation for automated remediation, and Amazon EventBridge for notification.
upvoted 3 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: C

Checks if the incoming SSH traffic for the security groups is accessible. The rule is COMPLIANT when IP addresses of the incoming SSH traffic in the security groups are restricted (CIDR other than 0.0.0.0/0). This rule applies only to IPv4.

Identifier: INCOMING_SSH_DISABLED

Resource Types: AWS::EC2::SecurityGroup

Trigger type: Configuration changes

upvoted 1 times

🗳️ 👤 **mana25** 1 year, 4 months ago

The solution also must remediate the security group rule automatically,

where that option remediate the issue?

upvoted 1 times

🗳️ 👤 **mana25** 1 year, 4 months ago

got it, with this part: DisablePublicAccessForSecurityGroup runbook

upvoted 1 times

🗨️ 👤 **tts1234** 2 years ago

Selected Answer: C

<https://docs.aws.amazon.com/config/latest/developerguide/restricted-ssh.html>

upvoted 4 times

🗨️ 👤 **jipark** 1 year, 4 months ago

"AWS Config" do this !

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years ago

Selected Answer: C

I go for C

upvoted 4 times

🗨️ 👤 **tyfta6** 2 years ago

Selected Answer: C

Vote for C

upvoted 3 times

A company has an application that runs only on Amazon EC2 Spot Instances. The instances run in an Amazon EC2 Auto Scaling group with scheduled scaling actions. However, the capacity does not always increase at the scheduled times, and instances terminate many times a day. A SysOps administrator must ensure that the instances launch on time and have fewer interruptions.

Which action will meet these requirements?

- A. Specify the capacity-optimized allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
- B. Specify the capacity-optimized allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.
- C. Specify the lowest-price allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
- D. Specify the lowest-price allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **r2c3po** 1 year ago

Selected Answer: A

To ensure that Spot Instances launch on time and experience fewer interruptions, you should use the capacity-optimized allocation strategy. This strategy helps you get the most capacity for your specified Spot requirements and reduces the likelihood of interruptions.

Additionally, adding more instance types to the Auto Scaling group enhances the chances of finding available Spot Instances at any given time. By diversifying the instance types, you increase the chances of securing Spot Instances with different pricing characteristics, which can contribute to a more stable and cost-effective configuration.

upvoted 2 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: A

The correct answer is A. Specify the capacity-optimized allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.

By using the capacity-optimized allocation strategy for Spot Instances, Amazon EC2 Auto Scaling will launch instances on the most available Spot Instance pools with the lowest prices. This helps to improve the chances of getting the required capacity at the scheduled times.

Adding more instance types to the Auto Scaling group also increases the chances of finding available Spot Instances at any given time. It provides flexibility in selecting different instance types based on the availability and cost of Spot Instances in different pools.

Option B (Specify the capacity-optimized allocation strategy for Spot Instances and increase the size of the instances in the Auto Scaling group) might not be the most efficient approach, as simply increasing the instance size may not necessarily address the issue of instances terminating frequently.

upvoted 2 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Option C (Specify the lowest-price allocation strategy for Spot Instances and add more instance types to the Auto Scaling group) does not align with the requirement of reducing interruptions since the lowest-price allocation strategy might not guarantee availability and can lead to more frequent interruptions.

Option D (Specify the lowest-price allocation strategy for Spot Instances and increase the size of the instances in the Auto Scaling group) could increase the chances of getting instances at a lower price, but it might not fully address the issue of instances not launching on time.

upvoted 1 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: A

aaaaaa

upvoted 2 times

🗳️ 👤 **Xelnak** 2 years, 1 month ago

Selected Answer: A

in August 2019 AWS launched the capacity-optimized allocation strategy for Spot Instances, which helps customers tap into the deepest Spot Instance pools by analyzing capacity metrics. Since then, customers have seen a significantly lower interruption rate with capacity-optimized strategy <https://aws.amazon.com/blogs/compute/introducing-price-capacity-optimized-allocation-strategy-for-ec2-spot-instances/>
upvoted 3 times

A company plans to deploy a database on an Amazon Aurora MySQL DB cluster. The database will store data for a demonstration environment. The data must be reset on a daily basis.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create a manual snapshot of the DB cluster after the data has been populated. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the snapshot and then delete the previous DB cluster.
- B. Enable the Backtrack feature during the creation of the DB cluster. Specify a target backtrack window of 48 hours. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to perform a backtrack operation.
- C. Export a manual snapshot of the DB cluster to an Amazon S3 bucket after the data has been populated. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the snapshot from Amazon S3.
- D. Set the DB cluster backup retention period to 2 days. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to restore the DB cluster to a point in time and then delete the previous DB cluster.

Suggested Answer: B

Community vote distribution

B (100%)

  **r2c3po** 1 year ago

Selected Answer: B

Option B provides an operationally efficient solution for resetting the data in an Amazon Aurora MySQL DB cluster on a daily basis:

B. Enable the Backtrack feature during the creation of the DB cluster. Specify a target backtrack window of 48 hours. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to perform a backtrack operation.
upvoted 1 times

  **Christina666** 1 year, 5 months ago



Selected Answer: B

The correct answer is B. Enable the Backtrack feature during the creation of the DB cluster. Specify a target backtrack window of 48 hours. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis. Configure the function to perform a backtrack operation.

Enabling the Backtrack feature in Amazon Aurora MySQL allows you to rewind the database cluster to a specific point in time without the need for manual snapshots. This feature is specifically designed for scenarios where you want to reset the data to a previous state, such as in a demonstration environment. By specifying a target backtrack window of 48 hours, you ensure that you have the ability to backtrack to a recent state.
upvoted 3 times

  **Christina666** 1 year, 5 months ago

Using an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function on a daily basis provides an automated and scheduled approach to resetting the data. The Lambda function can then perform the backtrack operation, making the process operationally efficient and requiring minimal manual intervention.
upvoted 2 times

  **Deeezz** 1 year, 11 months ago

The question says "The data must be reset on a daily basis." How can B be the answer when it states "Specify a target backtrack window of 48 hours." When did daily become 48 hours? is that a typo?
upvoted 2 times

  **Grodgar** 6 months, 1 week ago

MOST operationally optimal solution... A and C out (because manual never optimal.) D is not remembering the beginning state of the cluster. Only B left
upvoted 1 times

🗨️ 👤 **joanneli77** 1 year, 11 months ago

You can backtrack to any point in the last 48 hours, but will backtrack 24 hours. The extra day is arguable margin. (Actual Backtrack vs Target Backtrack)

upvoted 5 times

🗨️ 👤 **michaldavid** 2 years ago

Selected Answer: B

bbbbbbb

upvoted 2 times

🗨️ 👤 **tyfta6** 2 years ago

Selected Answer: B

Vote for B

upvoted 1 times

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

Ans: B

Ref: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.Backtrack.html>

upvoted 3 times

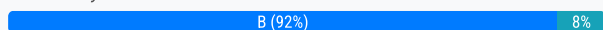
A SysOps administrator is setting up an automated process to recover an Amazon EC2 instance in the event of an underlying hardware failure. The recovered instance must have the same private IP address and the same Elastic IP address that the original instance had. The SysOps team must receive an email notification when the recovery process is initiated.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the `StatusCheckFailed_Instance` metric. Add an EC2 action to the alarm to recover the instance. Add an alarm notification to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the SysOps team email address to the SNS topic.
- B. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the `StatusCheckFailed_System` metric. Add an EC2 action to the alarm to recover the instance. Add an alarm notification to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the SysOps team email address to the SNS topic.
- C. Create an Auto Scaling group across three different subnets in the same Availability Zone with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP address and the Elastic IP address. Add an activity notification for the Auto Scaling group to send an email message to the SysOps team through Amazon Simple Email Service (Amazon SES).
- D. Create an Auto Scaling group across three Availability Zones with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP address and the Elastic IP address. Add an activity notification for the Auto Scaling group to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the SysOps team email address to the SNS topic.

Suggested Answer: D

Community vote distribution



tyfta6 Highly Voted 2 years ago

Selected Answer: B

Answer is B

You can create an Amazon CloudWatch alarm that monitors an Amazon EC2 instance and automatically recovers the instance if it becomes impaired due to an underlying hardware failure or a problem that requires AWS involvement to repair. Terminated instances cannot be recovered.

A recovered instance is identical to the original instance, including the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata. If the impaired instance has a public IPv4 address, the instance retains the public IPv4 address after recovery. If the impaired instance is in a placement group, the recovered instance runs in the placement group.

When the `StatusCheckFailed_System` alarm is triggered, and the recover action is initiated, you will be notified by the Amazon SNS topic that you selected when you created the alarm and associated the recover action.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

upvoted 8 times

james2033 Highly Voted 10 months, 1 week ago

Selected Answer: B

Option A like Option B so much, I read many time to see the different.

Read careful at <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-system-instance-status-check.html> 'StatusCheckFailed_System' and 'Hardware issues on the physical host that impact network reachability' --> Related to hardware --> Choose B.

Option A: 'StatusCheckFailed_Instance' related to kernal (Operation system software), networking, memory, file system.

upvoted 5 times

stoy123 Most Recent 10 months, 1 week ago

Selected Answer: B

ans. B

upvoted 2 times

🗨️ 👤 **r2c3po** 1 year ago

Selected Answer: D

By using an Auto Scaling group, you can achieve automated recovery in the event of an underlying hardware failure.

A launch template allows you to specify the private IP address and the Elastic IP address for the recovered instance, ensuring consistency with the original instance.

Activity notifications for the Auto Scaling group can be configured to publish messages to an Amazon SNS topic. Subscribing the SysOps team email address to the SNS topic ensures that the team receives email notifications when the recovery process is initiated.

upvoted 2 times

🗨️ 👤 **smanzana** 1 year, 5 months ago

Why not A???..... StatusCheckFailed_Instance metric focuses in hardware failures versus StatusCheckFailed_System metric focuses on issues at the operating system leveland the question says "in the event of a hardware failure"

upvoted 2 times

🗨️ 👤 **xSohox** 1 year, 4 months ago

No, the right answer is B.

instance check failure is related to Software or Networking issues on your instance

system check failure related to underlying hardware, etc.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-system-instance-status-check.html>

upvoted 2 times

🗨️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: B

System status checks monitor the AWS systems on which your instance runs. These checks detect underlying problems with your instance that require AWS involvement to repair. When a system status check fails, you can choose to wait for AWS to fix the issue, or you can resolve it yourself.

upvoted 3 times

🗨️ 👤 **Christina666** 1 year, 5 months ago

A recovered instance is identical to the original instance, including the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata. If the impaired instance has a public IPv4 address, the instance retains the public IPv4 address after recovery. If the impaired instance is in a placement group, the recovered instance runs in the placement group. During instance recovery, the instance is migrated as part of an instance reboot, and any data that is in-memory is lost.

upvoted 3 times

🗨️ 👤 **Christina666** 1 year, 5 months ago

Instance status checks monitor the software and network configuration of your individual instance. Amazon EC2 checks the health of the instance by sending an address resolution protocol (ARP) request to the network interface (NIC). These checks detect problems that require your involvement to repair. When an instance status check fails, you typically must address the problem yourself (for example, by rebooting the instance or by making instance configuration changes).

The following are examples of problems that can cause instance status checks to fail:

Failed system status checks

Incorrect networking or startup configuration

Exhausted memory

Corrupted file system

Incompatible kernel

If an instance status check fails, we increment the StatusCheckFailed_Instance metric.

upvoted 3 times

🗨️ 👤 **HSHan** 1 year, 6 months ago

Isn't this D?

I don't know what is the difference between A to B.

upvoted 1 times

🗨️ 👤 **maddy** 1 year, 5 months ago

For h/w failure - system check
For v/m or o/s failure - instance check
upvoted 2 times

  **guillaume01210** 1 year, 6 months ago



Selected Answer: B

where to ask the site to fix the preset answer ?
upvoted 2 times

  **michaldavid** 2 years ago

Selected Answer: B

bbbbbb
upvoted 2 times

  **Xelnak** 2 years, 1 month ago

Selected Answer: B

To automatically recover an instance when a system status check failure occurs, you can use the default configuration of the instance or create an Amazon CloudWatch alarm. If an instance becomes unreachable because of an underlying hardware failure or a problem that requires AWS involvement to repair, the instance is automatically recovered.

A recovered instance is identical to the original instance, including the instance ID, private IP addresses, Elastic IP addresses, and all instance metadata.

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

upvoted 4 times

A company has a public website that recently experienced problems. Some links led to missing webpages, and other links rendered incorrect webpages. The application infrastructure was running properly, and all the provisioned resources were healthy. Application logs and dashboards did not show any errors, and no monitoring alarms were raised. Systems administrators were not aware of any problems until end users reported the issues.

The company needs to proactively monitor the website for such issues in the future and must implement a solution as soon as possible.

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

- A. Rewrite the application to surface a custom error to the application log when issues occur. Automatically parse logs for errors. Create an Amazon CloudWatch alarm to provide alerts when issues are detected.
- B. Create an AWS Lambda function to test the website. Configure the Lambda function to emit an Amazon CloudWatch custom metric when errors are detected. Configure a CloudWatch alarm to provide alerts when issues are detected.
- C. Create an Amazon CloudWatch Synthetics canary. Use the CloudWatch Synthetics Recorder plugin to generate the script for the canary run. Configure the canary in line with requirements. Create an alarm to provide alerts when issues are detected.
- D. In the Amazon CloudWatch console, turn on Application Insights. Create a CloudWatch alarm to provide alerts when an issue is detected.

Suggested Answer: C

Community vote distribution

C (83%)

B (17%)

 **BugsBunny9998666** Highly Voted 2 years, 1 month ago

Selected Answer: C

C is possible, correct me if I'm wrong: Canaries are scripts written in Node.js or Python.

They create Lambda functions in your account that use Node.js or Python as a framework. Canaries work over both HTTP and HTTPS protocols, which makes it possible for you to continually verify your customer experience even when you don't have any customer traffic on your applications. By using canaries, you can discover issues before your customers do. This is why B is good but as C is automation of what B says making it a better option
upvoted 7 times

 **r2c3po** Most Recent 1 year ago

Selected Answer: C

Option C provides a solution with the least operational overhead for proactively monitoring the website:

Amazon CloudWatch Synthetics allows you to create canaries, which are scripts that mimic the behavior of a user interacting with your application.

The canary can be configured to navigate through the website, checking for missing or incorrect webpages.

The CloudWatch Synthetics Recorder plugin simplifies the process of creating canary scripts by recording user interactions with the website.

You can configure the canary to run at specific intervals to proactively monitor the website.

If issues are detected during canary runs, you can set up CloudWatch alarms to receive alerts.

upvoted 3 times

 **Christina666** 1 year, 5 months ago



Selected Answer: C

Option C: Create an Amazon CloudWatch Synthetics canary is the solution that will meet the requirements with the LEAST operational overhead.

Explanation:

CloudWatch Synthetics canaries are specifically designed for proactively monitoring websites and APIs. They simulate user interactions and monitor the website's functionality from the end-user's perspective. In this scenario, the website experienced problems with missing and incorrect webpages, which are issues that can be proactively monitored using a Synthetics canary.

upvoted 3 times



  **Christina666** 1 year, 5 months ago
Here's why Option C is the best choice:



CloudWatch Synthetics canaries are low operational overhead: They are managed by AWS and require minimal setup and maintenance. AWS takes care of the underlying infrastructure and scaling, so the operational burden is reduced for the company.



CloudWatch Synthetics Recorder plugin: The Synthetics Recorder plugin allows you to create the canary script easily without manual scripting. This saves time and effort for the administrators.



Monitoring from end-user perspective: Synthetics canaries simulate user interactions, ensuring that the website is functioning as expected for end-users. This is crucial in identifying issues that may not be apparent from the application logs or infrastructure metrics.



Proactive monitoring: Synthetics canaries continuously monitor the website, allowing for early detection of issues before end-users report them. This helps in addressing problems before they impact users and helps maintain the website's reliability.
upvoted 5 times



  **jipark** 1 year, 4 months ago
"CloudWatch Synthetics canaries" mimic "user behavior" and provide insights into the health and performance of your systems
upvoted 2 times



  **Christina666** 1 year, 5 months ago
Option A involves rewriting the application, which could be time-consuming and may introduce other complexities. Option B involves creating a custom Lambda function, which adds operational overhead in managing the function and ensuring its correctness. Option D involves turning on Application Insights, which might not be suitable for proactive monitoring of website issues and may not be as lightweight as CloudWatch Synthetics.
upvoted 2 times



  **dangji** 1 year, 11 months ago
Selected Answer: C
Canaries follow the same routes and perform the same actions as a customer
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Synthetics_Canaries.html
upvoted 2 times

  **michaldavid** 2 years ago
Selected Answer: C
cccccccc
upvoted 2 times

  **BugsBunny9998666** 2 years, 1 month ago
C is possible: Canaries are scripts written in Node.js or Python. They create Lambda functions in your account that use Node.js or Python as a framework. Canaries work over both HTTP and HTTPS protocols. which makes it possible for you to continually verify your customer experience even when you don't have any customer traffic on your applications. By using canaries, you can discover issues before your customers do.
upvoted 4 times

  **BugsBunny9998666** 2 years, 1 month ago
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch_Synthetics_Canaries.html
upvoted 3 times

  **Raynor** 2 years, 1 month ago
Selected Answer: B
This is B!
upvoted 3 times

  **zoltar_z** 2 years ago
C is the simplest way to test web sites
upvoted 2 times

A SysOps administrator is responsible for a company's security groups. The company wants to maintain a documented trail of any changes that are made to the security groups. The SysOps administrator must receive notification whenever the security groups change.

Which solution will meet these requirements?

- A. Set up Amazon Detective to record security group changes. Specify an Amazon CloudWatch Logs log group to store configuration history logs. Create an Amazon Simple Queue Service (Amazon SQS) queue for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SQS queue.
- B. Set up AWS Systems Manager Change Manager to record security group changes. Specify an Amazon CloudWatch Logs log group to store configuration history logs. Create an Amazon Simple Notification Service (Amazon SNS) topic for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SNS topic.
- C. Set up AWS Config to record security group changes. Specify an Amazon S3 bucket as the location for configuration snapshots and history files. Create an Amazon Simple Notification Service (Amazon SNS) topic for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SNS topic.
- D. Set up Amazon Detective to record security group changes. Specify an Amazon S3 bucket as the location for configuration snapshots and history files. Create an Amazon Simple Notification Service (Amazon SNS) topic for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SNS topic.

Suggested Answer: C

Community vote distribution

C (100%)

 **Raynor** Highly Voted 2 years, 1 month ago

Selected Answer: C

maintain a documented trail of any changes = config + S3

Notification = SNS

upvoted 6 times

 **Raynor** 2 years, 1 month ago

maintain a documented trail of any changes = config + S3,

Alert = SNS

upvoted 3 times

 **r2c3po** Most Recent 1 year ago

Selected Answer: C

Option C provides a solution using AWS Config to record security group changes and notify the SysOps administrator:

C. Set up AWS Config to record security group changes. Specify an Amazon S3 bucket as the location for configuration snapshots and history files. Create an Amazon Simple Notification Service (Amazon SNS) topic for notifications about configuration changes. Subscribe the SysOps administrator's email address to the SNS topic.

AWS Config can be configured to capture configuration changes, including changes to security groups.

The configuration snapshots and history files can be stored in an Amazon S3 bucket.

An Amazon SNS topic can be created to send notifications about configuration changes.

The SysOps administrator can subscribe their email address to the SNS topic to receive notifications whenever security group changes occur.

upvoted 1 times

 **Christina666** 1 year, 5 months ago

Selected Answer: C

AWS Config, a fully managed service that provides you with an AWS resource inventory, configuration history, and configuration change notifications to help enable security and governance. You can create AWS Config rules that automatically check the configuration of AWS resources that are recorded by AWS Config. For this example, I use a Config rule that is invoked whenever a change is made to a security group. Attach the Config rule to an AWS Lambda function that examines the ingress rules of a security group to see if the group remains in compliance with the rules.



upvoted 4 times

  **Christina666** 1 year, 5 months ago

The following Lambda function code defines a list named `REQUIRED_PERMISSIONS` with elements that represent a protocol, port range, and IP range that together define a security permission. This JSON notation is identical to what you would use when creating a security group with the AWS EC2 `authorize-security-group-ingress` command.

```
REQUIRED_PERMISSIONS = [  
  {  
    "IpProtocol" : "tcp",  
    "FromPort" : 80,  
    "ToPort" : 80,  
    "UserIdGroupPairs" : [],  
    "IpRanges" : [{"CidrIp" : "0.0.0.0/0"}],  
    "PrefixListIds" : []  
  },  
  {  
    "IpProtocol" : "tcp",  
    "FromPort" : 443,  
    "ToPort" : 443,  
    "UserIdGroupPairs" : [],  
    "IpRanges" : [{"CidrIp" : "0.0.0.0/0"}],  
    "PrefixListIds" : []  
  }  
]
```

upvoted 2 times

  **Gomer** 1 year, 8 months ago

Selected Answer: C

<https://aws.amazon.com/blogs/security/how-to-monitor-aws-account-configuration-changes-and-api-calls-to-amazon-ec2-security-groups/>

upvoted 2 times

An ecommerce company has built a web application that uses an Amazon Aurora DB cluster. The DB cluster includes memory optimized instance types with both a writer node and a reader node. Traffic volume changes throughout the day. During sudden traffic surges, Amazon CloudWatch metrics for the DB cluster indicate high RAM consumption and an increase in select latency.

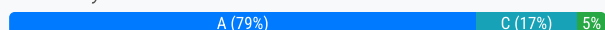
A SysOps administrator must implement a configuration change to improve the performance of the DB cluster. The change must minimize downtime and must not result in the loss of data.

Which change will meet these requirements?

- A. Add an Aurora Replica to the DB cluster.
- B. Modify the DB cluster to convert the DB cluster into a multi-master DB cluster.
- C. Take a snapshot of the DB cluster. From that snapshot, create a new DB cluster that has larger memory optimized instances.
- D. Increase the disk storage capacity of the DB cluster to double the existing disk capacity.

Suggested Answer: B

Community vote distribution



jipark Highly Voted 1 year, 10 months ago

Selected Answer: A

I'll give a tip :

read latency : read replica

write latency : multi-master

upvoted 12 times

Mila28 Highly Voted 2 years, 6 months ago

Selected Answer: A

The key word is "downtime", so I vote for A. B is out because this changes needs a downtime for business

upvoted 8 times

auxwww Most Recent 12 months ago

Selected Answer: A

'Select' latency - read-only workload causing high memory consumption and latency.

Therefore Read replica - A

upvoted 1 times

nyalpellymkar07 1 year, 1 month ago

Selected Answer: A

Keyword is select latency, which means read operations. So adding Read Replica will reduce Read Latency and also RAM consumption, load on the Primary.

upvoted 1 times

stoy123 1 year, 4 months ago

Selected Answer: A

ans. A

upvoted 2 times

r2c3po 1 year, 6 months ago

Selected Answer: A

A. Add an Aurora Replica to the DB cluster.

To improve the performance of the Amazon Aurora DB cluster with minimal downtime and no data loss, adding an Aurora Replica is a suitable option. Here's why:

Adding an Aurora Replica (Read Replica): This helps offload read traffic from the primary (writer) node to the replica (reader) node. It allows the read

workload to be distributed and can improve overall performance.

Minimal Downtime: Adding a read replica is a non-disruptive operation and can be performed with minimal downtime. It involves creating a read replica from the existing primary instance.

No Data Loss: Creating a read replica does not result in data loss. The replica is initially synchronized with the primary instance and continues to replicate changes as they occur.

--

Options B, C, and D involve more significant changes or potential downtime:

upvoted 2 times

🗳️ 👤 **Raj8989** 1 year, 8 months ago

Selected Answer: A

Keyword is downtime so it is A

upvoted 2 times

🗳️ 👤 **xile1021** 1 year, 8 months ago

Selected Answer: C

C

"If, after investigating your workload, you find that you need more memory, consider scaling up the DB instance class to a class with more RAM."

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

upvoted 2 times

🗳️ 👤 **wh1t4k3r** 1 year, 9 months ago

"The change must minimize downtime" <- its all about this quote

I agree that C would solve the issue in a better way, but migrating DBs is a operational nightmare.

upvoted 1 times

🗳️ 👤 **wh1t4k3r** 1 year, 9 months ago

A and B are valid. A vote A to be the simpler one. TO decide between both, more info is required.

upvoted 1 times

🗳️ 👤 **Christina666** 1 year, 11 months ago

Selected Answer: C

During sudden traffic surges, the high RAM consumption and increased select latency indicate that the current memory-optimized instance types might not be sufficient to handle the increased load. To improve the performance of the DB cluster, the best approach is to scale up the instance types to larger memory-optimized instances.

Option A (Add an Aurora Replica) would help with read scalability and high availability but may not directly address the high RAM consumption and select latency issues.

Option B (Modify the DB cluster to convert it into a multi-master DB cluster) doesn't directly address the high RAM consumption and select latency issues either. It would be more suitable for scenarios where you need to improve write scalability.

Option D (Increase the disk storage capacity) would only address storage-related issues and would not directly improve the performance of the DB cluster in terms of RAM consumption and select latency.

upvoted 4 times

🗳️ 👤 **Christina666** 1 year, 11 months ago

By taking a snapshot of the DB cluster and creating a new DB cluster with larger memory-optimized instances, you can effectively scale up the resources available to the database to handle the increased traffic volume and improve performance without losing any data or experiencing significant downtime. However, it's essential to plan and schedule this change during a maintenance window to minimize the impact on users.

upvoted 2 times

🗳️ 👤 **AShahine21** 1 year, 12 months ago

Not enough info to decide.

upvoted 1 times

🗳️ 👤 **noahsark** 2 years, 3 months ago

Selected Answer: A

Add an Aurora Replica to the DB cluster.

For select latency

upvoted 1 times

🗳️ 👤 **Vivec** 2 years, 3 months ago

Selected Answer: A

To improve the performance of the Aurora DB cluster, while minimizing downtime and without losing data, option A - adding an Aurora Replica to the DB cluster, is the best choice.

Adding an Aurora Replica would create a new reader node in the cluster, which can help distribute the load during sudden traffic surges, and reduce the select latency. This can also help reduce RAM consumption on the writer node. Additionally, adding a replica can be done with minimal downtime, as it does not require any changes to the application, and data will be automatically replicated from the writer node to the new replica.

upvoted 3 times

🗳️ 👤 **Gil80** 2 years, 4 months ago

Selected Answer: C

Option C, taking a snapshot of the DB cluster and creating a new DB cluster with larger memory-optimized instances, will allow the SysOps administrator to increase the memory capacity of the instances without losing any data. This option also minimizes downtime as the new DB cluster can be created from the snapshot while the original DB cluster continues to serve traffic. Once the new DB cluster is created, the administrator can update the web application to use the new DB cluster and then decommission the original DB cluster.

upvoted 2 times

🗳️ 👤 **defmania00** 2 years, 4 months ago

The DB cluster already includes memory optimized instances types and the issues show up during traffic surges and selects (READs). Using a replica should help with the selects.

upvoted 1 times

🗳️ 👤 **Gil80** 2 years, 4 months ago

Selected Answer: C

I say C.

The high RAM consumption and increased select latency suggest that the memory-optimized instances in the Aurora DB cluster are struggling to handle the sudden traffic surges. To improve the performance of the DB cluster, the SysOps administrator should increase the memory capacity of the instances.

Option A, adding an Aurora Replica, will not increase the memory capacity of the instances, so it is not a suitable option in this scenario.

Option B, modifying the DB cluster to convert it into a multi-master DB cluster, may improve write performance, but it is unlikely to improve read performance or memory capacity. Additionally, this option may result in downtime and require more significant configuration changes than other options.

Option D, increasing the disk storage capacity of the DB cluster, will not address the high RAM consumption and increased select latency, so it is not a suitable option in this scenario.

upvoted 3 times

🗳️ 👤 **zolthar_z** 2 years, 6 months ago

Selected Answer: A

Memory Issues and select latency are more related with read process

upvoted 2 times

🗳️ 👤 **tyfta6** 2 years, 6 months ago

Selected Answer: B

Vote for B

In a multi-master cluster, all DB instances have read/write capability. Multi-master clusters have different availability characteristics, support for database features, and procedures for monitoring and troubleshooting than single-master clusters.

upvoted 2 times

A company has a simple web application that runs on a set of Amazon EC2 instances behind an Elastic Load Balancer in the eu-west-2 Region. Amazon Route 53 holds a DNS record for the application with a simple routing policy. Users from all over the world access the application through their web browsers.

The company needs to create additional copies of the application in the us-east-1 Region and in the ap-south-1 Region. The company must direct users to the Region that provides the fastest response times when the users load the application.

What should a SysOps administrator do to meet these requirements?

- A. In each new Region, create a new Elastic Load Balancer and a new set of EC2 instances to run a copy of the application. Transition to a geolocation routing policy.
- B. In each new Region, create a copy of the application on new EC2 instances. Add these new EC2 instances to the Elastic Load Balancer in eu-west-2. Transition to a latency routing policy.
- C. In each new Region, create a copy of the application on new EC2 instances. Add these new EC2 instances to the Elastic Load Balancer in eu-west-2. Transition to a multivalue routing policy.
- D. In each new Region, create a new Elastic Load Balancer and a new set of EC2 instances to run a copy of the application. Transition to a latency routing policy.

Suggested Answer: D

Community vote distribution

D (100%)

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

Selected Answer: D

I'd say "latency routing" is the only solution to for "users from all over the world access the application" to have "fastest response times". If the users were limited to the three specified regions, then the answer might be A.

upvoted 6 times

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

Selected Answer: D

Option B was also pretty close in terms of transitioning over to the Latency routing selection IMO but the reason why I didn't pick it up was because it didn't specifically mention about the solution on the missing "ap-south-1" region. That missing gap in the statement leans me towards the option D.

upvoted 1 times

  **tyfta6** 2 years, 6 months ago

Does Elastic load balancer work across regions?

Amazon has made the creation and management of load balancers in the cloud a lot simpler when they created elastic load balancers. But elastic load balancers have one fatal flaw. They exist within a single AWS region and if that region is having an outage, then your whole application goes down.

upvoted 2 times

  **jipark** 1 year, 10 months ago

latency routing + ELB on one region seems clue.



upvoted 1 times

  **tyfta6** 2 years, 6 months ago

Selected Answer: D

Vote for D. Anybody look into B choice?

upvoted 2 times

  **beznika** 2 years, 6 months ago



Correct is D. The fastest response time so latency based routing has to be used.

upvoted 2 times

  **grka25** 2 years, 6 months ago

Answer is A

upvoted 1 times

  **grka25** 2 years, 6 months ago

Correction, D

upvoted 1 times

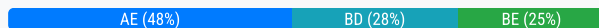
A company creates a new member account by using AWS Organizations. A SysOps administrator needs to add AWS Business Support to the new account.

Which combination of steps must the SysOps administrator take to meet this requirement? (Choose two.)

- A. Sign in to the new account by using IAM credentials. Change the support plan.
- B. Sign in to the new account by using root user credentials. Change the support plan.
- C. Use the AWS Support API to change the support plan.
- D. Reset the password of the account root user.
- E. Create an IAM user that has administrator privileges in the new account.

Suggested Answer: *BD*

Community vote distribution



  **Untamables** Highly Voted  2 years, 4 months ago

Selected Answer: AE

Correct answers are A and E.

The combination of option B and D also works, but the best practice is using an IAM user, not the root user.

<https://docs.aws.amazon.com/awssupport/latest/user/changing-support-plans.html>

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html

upvoted 12 times

 Christina666 1 year, 11 months ago

<https://docs.aws.amazon.com/awssupport/latest/user/changing-support->

plans.html#:~:text=To%20change%20your%20support%20plan%2C%20you%20must%20have%20AWS%20Identity%20and%20Access%20Management%20

upvoted 4 times

 eboehm Highly Voted 1 year, 11 months ago

Selected Answer: BD

I think its B and D. When a new account joins organizations, it creates a new temp root password and you will be prompted to change it when logging in for the first time.

"When you create a new account, AWS Organizations initially assigns a password to the root user that is a minimum of 64 characters long. All characters are randomly generated with no guarantees on the appearance of certain character sets. You can't retrieve this initial password. To access the account as the root user for the first time, you must go through the process for password recovery."

upvoted 7 times

 Rado_Piatek Most Recent 1 month ago

Selected Answer: BD

B and D: to change Support plan you need to be logged in as root

upvoted 1 times

 Rado_Piatek 1 month, 1 week ago

Selected Answer: BD

B: Only the root user of the AWS account can change the support plan.

D: It is a new account, so it will be necessary to change it.

upvoted 1 times

 numark 6 months, 1 week ago

Selected Answer: BD

Ok after extensive research and running this through three different AIs it is definitely B & D. To add AWS Business Support to the new account, You HAVE to be logged in as the root user. Admin will not work.

upvoted 1 times

 0c2d840 6 months, 1 week ago



Selected Answer: AD

Answer is A and D. When using Organizations, use of root account to manager member accounts is not recommended. It should be accessed using non-root IAM user having appropriate permissions.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html

"New accounts you create in Organizations have no root user credentials by default."

upvoted 1 times

  **0c2d840** 6 months, 1 week ago



Correction:

Answer is A and E. When using Organizations, use of root account to manager member accounts is not recommended. It should be accessed using non-root IAM user having appropriate permissions.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html



"New accounts you create in Organizations have no root user credentials by default."

upvoted 1 times

  **numark** 7 months, 1 week ago

how can you pick D: This step is unnecessary unless the root user credentials are lost. It is not a prerequisite to upgrading the support plan.



upvoted 1 times

  **numark** 7 months, 1 week ago

A & B: Changing the AWS Support plan requires administrative-level access. The root user has unrestricted access to manage account-level settings, including billing and support plans.

AWS recommends using the root user credentials only when performing account setup tasks like upgrading the support plan. If an IAM user with the appropriate "aws-portal:ModifyAccount" or billing permissions exists, they can also upgrade the support plan. This approach aligns with the principle of least privilege by avoiding root user usage whenever possible. If an administrator IAM user already exists, they can perform the necessary actions. Creating a new IAM user is unrelated to the support plan upgrade process.

upvoted 1 times

  **Aamee** 9 months, 1 week ago

Selected Answer: AE

First, let's have the Admin user created via "option E" and then let's get signed in via that newly admin user and have the support plan changed by that admin user. Hence, this combination follows one of the security best practices by performing the operations via least privileged access.

upvoted 1 times

  **tgiv** 1 year, 2 months ago

Selected Answer: BD

I will go with BD as well. Upon my research, you don't need administrator access on a IAM user in order to change the support plan, so A and E are out of question because you are ignoring the least privilege principle.

Regarding the AWS Support API, this is used to open cases and other actions related to support and NOT to change the support plan.

So in this scenario it would be the best practice to login as root to change the support plan and change the password.

upvoted 1 times

  **March2023** 1 year, 3 months ago

Selected Answer: AE

going with A and E



upvoted 2 times

  **Rabbit117** 1 year, 3 months ago

Selected Answer: AE

An IAM user with correct permissions or the root user can change support plans, so option B would work. However, the question says to choose a combination of steps, 2 answers. Therefore I think E and A are correct.



upvoted 2 times

  **stoy123** 1 year, 4 months ago

Selected Answer: AE

A and E

upvoted 2 times

  **r2c3po** 1 year, 6 months ago

Selected Answer: BE

To add AWS Business Support to a new member account in AWS Organizations, you need to perform the following steps:

Sign in as the Root User:

Sign in to the new account using the root user credentials.

Change Support Plan:

After signing in as the root user, navigate to the AWS Support Center.

Change the support plan to AWS Business Support.

Create IAM User (Optional):

Optionally, to enhance security and follow best practices, create an IAM user with administrator privileges.

Use the IAM user for day-to-day operations instead of relying on the root user.

Options A and C are incorrect:

upvoted 1 times

  **Hatem08** 1 year, 6 months ago

Selected Answer: AE



create an IAM user that has administrator privileges in the new account.

Sign in to the new account by using IAM credentials. Change the support plan.

<https://docs.aws.amazon.com/awssupport/latest/user/changing-support-plans.html>

<https://docs.aws.amazon.com/awssupport/latest/user/security-support-plans.html>

upvoted 4 times

  **Mila28** 1 year, 7 months ago

Selected Answer: BD

I'll give some tips:

A: you can't change plan with basic user

C: Support plan not has an API

E: Not is necessary for change support plan

So you need access as root user, maybe if you don't have the password you maybe need reset password

upvoted 5 times

  **callspace** 1 year, 8 months ago

Selected Answer: BE

As per the link from @ Christina666

To change your support plan, you must have AWS Identity and Access Management (IAM) permissions or sign in to your account as a root user..

upvoted 3 times

A SysOps administrator creates two VPCs, VPC1 and VPC2, in a company's AWS account. The SysOps administrator deploys a Linux Amazon EC2 instance in VPC1 and deploys an Amazon RDS for MySQL DB instance in VPC2. The DB instance is deployed in a private subnet. An application that runs on the EC2 instance needs to connect to the database.

What should the SysOps administrator do to give the EC2 instance the ability to connect to the database?

- A. Enter the DB instance connection string into the VPC1 route table.
- B. Configure VPC peering between the two VPCs.
- C. Add the same IPv4 CIDR range for both VPCs.
- D. Connect to the DB instance by using the DB instance's public IP address.

Suggested Answer: B

Community vote distribution

B (100%)

  **michaldavid** Highly Voted 1 year ago

Selected Answer: B

bbbbbbb

upvoted 5 times

  **bamishr** 1 year ago

hi michal when you are giving the exam?

upvoted 1 times

  **Fatoch** Most Recent 1 year ago

B is correct

upvoted 1 times

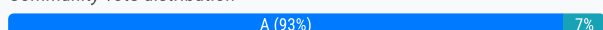
A company uses an Amazon S3 bucket to store data files. The S3 bucket contains hundreds of objects. The company needs to replace a tag on all the objects in the S3 bucket with another tag.

What is the MOST operationally efficient way to meet this requirement?

- A. Use S3 Batch Operations. Specify the operation to replace all object tags.
- B. Use the AWS CLI to get the tags for each object. Save the tags in a list. Use S3 Batch Operations. Specify the operation to delete all object tags. Use the AWS CLI and the list to retag the objects.
- C. Use the AWS CLI to get the tags for each object. Save the tags in a list. Use the AWS CLI and the list to remove the object tags. Use the AWS CLI and the list to retag the objects.
- D. Use the AWS CLI to copy the objects to another S3 bucket. Add the new tag to the copied objects. Delete the original objects.

Suggested Answer: A

Community vote distribution



Liongeek Highly Voted 2 years, 1 month ago

Ans: A

Ref. <https://aws.amazon.com/es/blogs/storage/adding-and-removing-object-tags-with-s3-batch-operations/>

upvoted 16 times

Gomer Highly Voted 1 year, 8 months ago

Selected Answer: A

"With this feature, you can make changes to object metadata and properties, or perform other storage management tasks, such as "replacing object tag sets"

<https://aws.amazon.com/s3/features/batch-operations/>

upvoted 5 times

r2c3po Most Recent 1 year ago

Selected Answer: A

S3 Batch Operations is the most operationally efficient way to meet the requirement of replacing a tag on all objects in an S3 bucket. S3 Batch Operations allows you to perform large-scale batch operations on Amazon S3 objects, including updating tags on objects.

Here's why option A is the best choice:

S3 Batch Operations:

Allows you to perform actions, such as tagging, on a large number of objects at once.

Is designed for efficiency and scalability when dealing with large-scale operations.

Is a native AWS service specifically built for performing batch operations on S3 objects.

upvoted 4 times

csG13 1 year, 9 months ago

Selected Answer: A

A is by far the MOST operationally efficient way, see docs below

<https://aws.amazon.com/blogs/storage/adding-and-removing-object-tags-with-s3-batch-operations/>

upvoted 4 times

Gil80 1 year, 10 months ago

Selected Answer: A

A is the correct answer.

The question is asking to replace existing tag with another, not adding new tags.

Therefore, batch operations is correct (A).

Source: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/batch-ops-put-object-tagging.html>

upvoted 5 times

🗨️ 👤 **joanneli77** 1 year, 11 months ago

Selected Answer: C

If you replace all object tags, existing tags that you do not want to replace will be lost. You'll need to read those tags to retain them and update the tags you're trying to update.

upvoted 1 times

🗨️ 👤 **Gil80** 1 year, 10 months ago

The question didn't specify that there are other tags that shouldn't be erased. It did mention the MOST operationally efficient way, which is A not C because they specify it has hundreds of objects and C is not efficient for large buckets.

A seems to be the better answer, even if C is doable.

upvoted 1 times

🗨️ 👤 **defmania00** 1 year, 10 months ago

Interesting part with C is the fact that nowhere in that answer you're using the new tags. You're getting the existing tags, saving them in a list, use CLI to remove the tags and then use that same list to retag. Won't that get you the exact same tags you had in the first place?

upvoted 1 times

A company needs to take an inventory of applications that are running on multiple Amazon EC2 instances. The company has configured users and roles with the appropriate permissions for AWS Systems Manager. An updated version of Systems Manager Agent has been installed and is running on every instance. While configuring an inventory collection, a SysOps administrator discovers that not all the instances in a single subnet are managed by Systems Manager.

What must the SysOps administrator do to fix this issue?

- A. Ensure that all the EC2 instances have the correct tags for Systems Manager access.
- B. Configure AWS Identity and Access Management Access Analyzer to determine and automatically remediate the issue.
- C. Ensure that all the EC2 instances have an instance profile with Systems Manager access.
- D. Configure Systems Manager to use an interface VPC endpoint.

Suggested Answer: D

Community vote distribution



Liongeek Highly Voted 2 years, 1 month ago

It looks like C for me. I don't think it's D. In order to be D, all instances in that subnet shouldn't be in SSM but question says that SOME INSTANCES IN THE SAME SUBNET...

upvoted 9 times

r2c3po Most Recent 1 year ago

Selected Answer: C

Here's why option C is the correct choice:

Ensure IAM Instance Profile:

Systems Manager relies on IAM roles and instance profiles to access and manage EC2 instances.

The IAM instance profile associated with an EC2 instance must have the necessary permissions to allow Systems Manager to perform operations on the instance.

upvoted 4 times

Mila28 1 year, 1 month ago

Selected Answer: C

Some tips for choose correct answer:

A: Is for map tags not for fix some issue with ssm

B: Is for review rules or configs not for issue connect with ssm

D: With this you can comunicate ssm with ec2, but you can't fix the issue connectivity

So C is the best option

upvoted 3 times

pepecastr0 1 year, 6 months ago

For sure is C

upvoted 1 times

tinyflame 1 year, 8 months ago

Selected Answer: D

Network Problem

upvoted 1 times

tinyflame 1 year, 8 months ago

Sorry D is incorrect.

Because VPC Endpoint Configuration is for Subnet not for Systems Manager.

upvoted 2 times

michaldavid 2 years ago

Selected Answer: C

cccccc

upvoted 3 times

  **tyfta6** 2 years ago

Selected Answer: C

Vote for C

upvoted 3 times

  **BugsBunny9998666** 2 years ago

Selected Answer: C

C makes most cense

upvoted 4 times

  **CloudHandsOn** 2 years ago

I'll go with C as well

upvoted 3 times

A company stores sensitive data in an Amazon S3 bucket. The company must log all access attempts to the S3 bucket. The company's risk team must receive immediate notification about any delete events.

Which solution will meet these requirements?

- A. Enable S3 server access logging for audit logs. Set up an Amazon Simple Notification Service (Amazon SNS) notification for the S3 bucket. Select DeleteObject for the event type for the alert system.
- B. Enable S3 server access logging for audit logs. Launch an Amazon EC2 instance for the alert system. Run a cron job on the EC2 instance to download the access logs each day and to scan for a DeleteObject event.
- C. Use Amazon CloudWatch Logs for audit logs. Use Amazon CloudWatch alarms with an Amazon Simple Notification Service (Amazon SNS) notification for the alert system.
- D. Use Amazon CloudWatch Logs for audit logs. Launch an Amazon EC2 instance for the alert system. Run a cron job on the EC2 instance each day to compare the list of the items with the list from the previous day. Configure the cron job to send a notification if an item is missing.

Suggested Answer: A

Community vote distribution

A (90%)

10%

 **Granddude** 3 months, 2 weeks ago

Selected Answer: C

Although S3 server access logging (as mentioned in Option A) captures all access events, its delivery is best-effort and often delayed—making it unsuitable for immediate notifications. By using CloudTrail to log S3 data events to CloudWatch Logs and creating metric filters and alarms for DeleteObject events, you can trigger near real-time SNS notifications to the risk team. This solution meets both the requirement for comprehensive logging and immediate alerts.

AWS Documentation: Logging Data Events with CloudTrail and Using Amazon CloudWatch Alarms

upvoted 1 times


 **XXXXXINN** 7 months, 2 weeks ago

I think the key here is asking for 'Immediate notifications...'.

The list of options suck. for A, SNS has to worked with S3 Event Notifications in order to provide real time trigger and notifications. For C, CloudWatch alone without CloudTrail deliveries logs to it, then it useless...

so I have no idea which one is the official answer - my bet here is this question will never appear in exam anymore - it's disqualified :)


upvoted 1 times

 **alexleely** 1 year, 10 months ago

Selected Answer: A

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html>

upvoted 2 times

 **Gomer** 2 years, 2 months ago

Selected Answer: A

Answer should be S3 Data logging with CloudTrail. However, that is not listed as an option. Next best answer is A. EC2 answers are nonsensical . CloudWatch Logs is not is not CloudTrail data logging.

upvoted 2 times

 **michele_scar** 2 years, 3 months ago

Selected Answer: A

as the doc explain

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/notification-how-to-event-types-and-destinations.html#supported-notification-event-types>

You should trigger object action to SNS, SQS and Lambda

upvoted 1 times

 **defmania00** 2 years, 4 months ago

The answers seem incomplete. Enabling S3 server access logging is very good for performing security and access audits. However, these logs are stored in an S3 bucket, when enabled. As for the delete events, wouldn't Event Notifications be the way to go? Create an event notification for object removal and use an SNS topic to send the notifications to your risk team.

upvoted 1 times

🗨️ 👤 **Untamables** 2 years, 4 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging.html>

Option A and B are wrong. The completeness and timeliness of server access logging is not guaranteed.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerLogs.html#LogDeliveryBestEffort>

upvoted 2 times

🗨️ 👤 **defmania00** 2 years, 4 months ago

C is not CloudTrail, C is talking about CloudWatch logs.

upvoted 1 times

🗨️ 👤 **MrMLB** 2 years, 6 months ago

Selected Answer: A

A

To meet the requirements of logging all access attempts to the S3 bucket and receiving immediate notification about any delete events, the company can enable S3 server access logging and set up an Amazon Simple Notification Service (Amazon SNS) notification for the S3 bucket. The S3 server access logs will record all access attempts to the bucket, including delete events, and the SNS notification can be configured to send an alert when a DeleteObject event occurs.

upvoted 3 times

🗨️ 👤 **beznika** 2 years, 6 months ago

A looks good.

upvoted 1 times

🗨️ 👤 **michaldavid** 2 years, 6 months ago

Selected Answer: A

aaaaaaa

upvoted 2 times

🗨️ 👤 **Fatoch** 2 years, 6 months ago

is it not C?

upvoted 1 times

A SysOps administrator receives an alert from Amazon GuardDuty about suspicious network activity on an Amazon EC2 instance. The GuardDuty finding lists a new external IP address as a traffic destination. The SysOps administrator does not recognize the external IP address. The SysOps administrator must block traffic to the external IP address that GuardDuty identified.

Which solution will meet this requirement?

- A. Create a new security group to block traffic to the external IP address. Assign the new security group to the EC2 instance.
- B. Use VPC flow logs with Amazon Athena to block traffic to the external IP address.
- C. Create a network ACL. Add an outbound deny rule for traffic to the external IP address.
- D. Create a new security group to block traffic to the external IP address. Assign the new security group to the entire VPC.

Suggested Answer: C

Community vote distribution

C (100%)

  **beznika** Highly Voted 2 years, 6 months ago

Security groups are out because you allow traffic using security groups not block. VPC flow logs with Athena? How can that help? And the ACL outbound rule to block the IP? ACL makes the most sense because if the IP is the destination the outbound rule to block will do. However it would make more sense to modify existing ACL because a subnet can be associated with only one ACL. So I am going to say C is the correct one.

upvoted 6 times

  **10cc6ba** Most Recent 11 months, 3 weeks ago

Selected Answer: C

C only 100%

upvoted 2 times

  **Christina666** 1 year, 11 months ago

Selected Answer: C

C. Create a network ACL. Add an outbound deny rule for traffic to the external IP address.

Explanation: Network Access Control Lists (NACLs) are used to control the traffic entering and exiting subnets in a VPC. They operate at the subnet level and are stateless, meaning that both inbound and outbound rules must be explicitly defined. By adding an outbound deny rule for traffic to the specific external IP address identified by GuardDuty, you can block any communication from the EC2 instance to that IP address.

upvoted 1 times


  **Christina666** 1 year, 11 months ago

Option A (Create a new security group to block traffic to the external IP address and assign it to the EC2 instance) is incorrect because security groups control inbound and outbound traffic to and from an EC2 instance, but they cannot be used to block traffic to external IP addresses. They only allow you to specify allowed traffic based on ports and protocols.

Option B (Use VPC flow logs with Amazon Athena to block traffic to the external IP address) is incorrect because VPC flow logs do not have the capability to block traffic. They are used for monitoring and logging network traffic, but they cannot be used for active traffic control.

Option D (Create a new security group to block traffic to the external IP address and assign it to the entire VPC) is incorrect because, like in Option A, security groups are not meant to block traffic to specific external IP addresses. Assigning the security group to the entire VPC will not achieve the goal of blocking traffic to the specific IP address identified by GuardDuty.

upvoted 2 times

  **jipark** 1 year, 10 months ago

why not A. security group can only allow traffic.


why C. NACL deny/allow traffic by creating new one.

upvoted 2 times

  **pepecastr0** 2 years ago

C, best way to block outbound traffic, but I'm not sure why you need to create a new NACL instead of add the rule to the existing one

upvoted 1 times

  **zolthar_z** 2 years, 6 months ago

Selected Answer: C

Answer is C, ACL is the only way to block outbound traffic

upvoted 2 times

A company's reporting job that used to run in 15 minutes is now taking an hour to run. An application generates the reports. The application runs on Amazon EC2 instances and extracts data from an Amazon RDS for MySQL database.

A SysOps administrator checks the Amazon CloudWatch dashboard for the RDS instance and notices that the Read IOPS metrics are high, even when the reports are not running. The SysOps administrator needs to improve the performance and the availability of the RDS instance.

Which solution will meet these requirements?

- A. Configure an Amazon ElastiCache cluster in front of the RDS instance. Update the reporting job to query the ElastiCache cluster.
- B. Deploy an RDS read replica. Update the reporting job to query the reader endpoint.
- C. Create an Amazon CloudFront distribution. Set the RDS instance as the origin. Update the reporting job to query the CloudFront distribution.
- D. Increase the size of the RDS instance.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **zolthar_z** Highly Voted 👍 1 year, 6 months ago

Selected Answer: B

If you have read issues read-replica is always the solution
upvoted 7 times

🗳️ 👤 **Salah94** Most Recent 🕒 5 months, 1 week ago

Selected Answer: A

Because it says "even when the reports are not running", so I would go with A as adding cache level shall eliminate read IOPS load on the database.
upvoted 1 times

🗳️ 👤 **Christina666** 11 months, 1 week ago

Selected Answer: B

BBBBB.....
upvoted 3 times

🗳️ 👤 **joanneli77** 1 year, 4 months ago

There is not enough information to make a distinction between A and B. Separating I/O from writes and reads, vs. caching, depends on how write intensive the RDS is. If it's doing nearly zero writes, a read replica at the same I/O level won't solve the problem. That information is not provided. A or B, depending on data.
upvoted 2 times

🗳️ 👤 **defmania00** 1 year, 4 months ago

I think the keyword is improve availability. It does seem that B is the right answer.
upvoted 6 times

🗳️ 👤 **beznika** 1 year, 6 months ago

My choice B
upvoted 3 times

A company's SysOps administrator regularly checks the AWS Personal Health Dashboard in each of the company's accounts. The accounts are part of an organization in AWS Organizations. The company recently added 10 more accounts to the organization. The SysOps administrator must consolidate the alerts from each account's Personal Health Dashboard.

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

- A. Enable organizational view in AWS Health.
- B. Configure the Personal Health Dashboard in each account to forward events to a central AWS CloudTrail log.
- C. Create an AWS Lambda function to query the AWS Health API and to write all events to an Amazon DynamoDB table.
- D. Use the AWS Health API to write events to an Amazon DynamoDB table.

Suggested Answer: A

Community vote distribution

A (100%)

 **zolthar_z** Highly Voted 2 years ago

Selected Answer: A

The complete answer is: Go to Admin account and add the health dashboards from other accounts
upvoted 7 times

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

Selected Answer: A

<https://docs.aws.amazon.com/health/latest/ug/enable-organizational-view-in-health-console.html>
upvoted 3 times

 **Christina666** 1 year, 5 months ago

Selected Answer: A

A. Enable organizational view in AWS Health.

Explanation: By enabling organizational view in AWS Health, you can aggregate and view the status of AWS Personal Health Dashboard across all accounts in your organization in AWS Organizations. This feature provides a centralized view of the health of all accounts, making it easier for the SysOps administrator to monitor and respond to alerts from a single dashboard rather than checking each account individually.

upvoted 2 times

A company runs an application on Amazon EC2 instances. The EC2 instances are in an Auto Scaling group and run behind an Application Load Balancer (ALB). The application experiences errors when total requests exceed 100 requests per second. A SysOps administrator must collect information about total requests for a 2-week period to determine when requests exceeded this threshold.

What should the SysOps administrator do to collect this data?

- A. Use the ALB's RequestCount metric. Configure a time range of 2 weeks and a period of 1 minute. Examine the chart to determine peak traffic times and volumes.
- B. Use Amazon CloudWatch metric math to generate a sum of request counts for all the EC2 instances over a 2-week period. Sort by a 1-minute interval.
- C. Create Amazon CloudWatch custom metrics on the EC2 launch configuration templates to create aggregated request metrics across all the EC2 instances.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Configure an EC2 event matching pattern that creates a metric that is based on EC2 requests. Display the data in a graph.

Suggested Answer: A

Community vote distribution

A (100%)

 **Christina666** Highly Voted 11 months, 1 week ago

Selected Answer: A

A. Use the ALB's RequestCount metric. Configure a time range of 2 weeks and a period of 1 minute. Examine the chart to determine peak traffic times and volumes.

Explanation: The ALB's RequestCount metric provides the total number of requests processed by the ALB. By configuring a time range of 2 weeks with a 1-minute period, you can collect detailed data on request counts for each minute over the 2-week period. This will allow the SysOps administrator to visualize the traffic patterns and identify peak times when the requests exceeded the threshold of 100 requests per second.

upvoted 6 times

 **zolthar_z** Most Recent 1 year, 6 months ago

Selected Answer: A

Answer is A

upvoted 2 times

A company recently migrated its application to a VPC on AWS. An AWS Site-to-Site VPN connection connects the company's on-premises network to the VPC. The application retrieves customer data from another system that resides on premises. The application uses an on-premises DNS server to resolve domain records. After the migration, the application is not able to connect to the customer data because of name resolution errors.

Which solution will give the application the ability to resolve the internal domain names?

- A. Launch EC2 instances in the VPC. On the EC2 instances, deploy a custom DNS forwarder that forwards all DNS requests to the on-premises DNS server. Create an Amazon Route 53 private hosted zone that uses the EC2 instances for name servers.
- B. Create an Amazon Route 53 Resolver outbound endpoint. Configure the outbound endpoint to forward DNS queries against the on-premises domain to the on-premises DNS server.
- C. Set up two AWS Direct Connect connections between the AWS environment and the on-premises network. Set up a link aggregation group (LAG) that includes the two connections. Change the VPC resolver address to point to the on-premises DNS server.
- D. Create an Amazon Route 53 public hosted zone for the on-premises domain. Configure the network ACLs to forward DNS requests against the on-premises domain to the Route 53 public hosted zone.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **Christina666** 11 months, 1 week ago

Selected Answer: B

outbound resolver: resolver on-prem connection

Inbound resolver: resolver dns in vpc

upvoted 4 times

🗳️ 👤 **dangji** 1 year, 5 months ago

Selected Answer: B

To forward DNS queries that originate on Amazon EC2 instances in one or more VPCs to your network

https://docs.aws.amazon.com/zh_tw/Route53/latest/DeveloperGuide/resolver-forwarding-outbound-queries.html

upvoted 3 times

🗳️ 👤 **michaldavid** 1 year, 6 months ago

Selected Answer: B

bbbbbb

upvoted 2 times

A company's web application is available through an Amazon CloudFront distribution and directly through an internet-facing Application Load Balancer (ALB). A SysOps administrator must make the application accessible only through the CloudFront distribution and not directly through the ALB. The SysOps administrator must make this change without changing the application code.

Which solution will meet these requirements?

- A. Modify the ALB type to internal. Set the distribution's origin to the internal ALB domain name.
- B. Create a Lambda@Edge function. Configure the function to compare a custom header value in the request with a stored password and to forward the request to the origin in case of a match. Associate the function with the distribution.
- C. Replace the ALB with a new internal ALB. Set the distribution's origin to the internal ALB domain name. Add a custom HTTP header to the origin settings for the distribution. In the ALB listener, add a rule to forward requests that contain the matching custom header and the header's value. Add a default rule to return a fixed response code of 403.
- D. Add a custom HTTP header to the origin settings for the distribution. In the ALB listener, add a rule to forward requests that contain the matching custom header and the header's value. Add a default rule to return a fixed response code of 403.

Suggested Answer: A

Community vote distribution

D (95%)

5%

 **Arnaud92** Highly Voted 2 years ago

Selected Answer: D

D

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

upvoted 5 times

 **johnson_chao** Most Recent 9 months ago

Selected Answer: D

answer is D

Configure CloudFront to add a custom HTTP header to requests that it sends to the Application Load Balancer.

ref : <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

upvoted 2 times

 **r2c3po** 1 year ago

Selected Answer: C

Option C provides a solution to make the application accessible only through the CloudFront distribution and not directly through the ALB.

upvoted 1 times

 **TwinSpark** 1 year, 2 months ago

Selected Answer: D

ALB cannot be internal.

<https://repost.aws/questions/QUdr1a-rXWQFiGqKVjNPhOpg/cloudfront-internal-elb-origin>

so it's D

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

upvoted 2 times

 **konieczny69** 1 year ago

Nonsense - <https://docs.aws.amazon.com/AmazonECS/latest/developerguide/create-application-load-balancer.html>

upvoted 1 times

 **Madiba237** 1 year, 3 months ago

The answer is C instead of D , D do not block access directly from the ALB

upvoted 2 times

 **konieczny69** 1 year ago

of course its C

upvoted 1 times

🗨️ 👤 **jipark** 1 year, 4 months ago

Selected Answer: D

it's similar what we do for SSL redirecting.
when came in directly, redirect L4 or DNS.

upvoted 1 times

🗨️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: D

All requests go through CCloudfront, sample cloudformation template as below:

DistributionConfig:

Origins:

- DomainName: app-load-balancer.example.com

Id: Example-ALB

CustomOriginConfig:

OriginProtocolPolicy: https-only

OriginSSLProtocols:

- TLSv1.2

OriginCustomHeaders:

- HeaderName: X-Custom-Header

HeaderValue: random-value-1234567890

upvoted 2 times

🗨️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: D

You can configure CloudFront to add a custom HTTP header to the requests that it sends to your origin (in this case, an Application Load Balancer).

Important

This use case relies on keeping the custom header name and value secret. If the header name and value are not secret, other HTTP clients could potentially include them in requests that they send directly to the Application Load Balancer. This can cause the Application Load Balancer to behave as though the requests came from CloudFront when they did not. To prevent this, keep the custom header name and value secret.

upvoted 2 times

🗨️ 👤 **Christina666** 1 year, 5 months ago

CloudFront can also help to reduce latency and even absorb some distributed denial of service (DDoS) attacks. However, if users can bypass CloudFront and access your Application Load Balancer directly, you don't get these benefits. But you can configure Amazon CloudFront and your Application Load Balancer to prevent users from directly accessing the Application Load Balancer. This allows users to access the Application Load Balancer only through CloudFront, ensuring that you get the benefits of using CloudFront.

To prevent users from directly accessing an Application Load Balancer and allow access only through CloudFront, complete these high-level steps:

Configure CloudFront to add a custom HTTP header to requests that it sends to the Application Load Balancer.

Configure the Application Load Balancer to only forward requests that contain the custom HTTP header.

(Optional) Require HTTPS to improve the security of this solution.

upvoted 2 times

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

Selected Answer: D

Add a custom HTTP header to the origin settings for the distribution. In the ALB listener, add a rule to forward requests that contain the matching custom header and the header's value. Add a default rule to return a fixed response code of 403.

Configure CloudFront to add a custom HTTP header to requests that it sends to the Application Load Balancer.

Configure the Application Load Balancer to only forward requests that contain the custom HTTP header.

(Optional) Require HTTPS to improve the security of this solution.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

upvoted 3 times

🗨️ 👤 **MrMLB** 2 years ago

Selected Answer: D

D

To make the application accessible only through the CloudFront distribution and not directly through the Application Load Balancer (ALB), you can add a custom HTTP header to the origin settings for the CloudFront distribution. You can then create a rule in the ALB listener to forward requests that contain the matching custom header and its value to the origin. You can also add a default rule to the ALB listener to return a fixed response code of 403 for requests that do not contain the matching custom header. This will allow you to redirect all requests to the CloudFront distribution and block direct access to the application through the ALB.

upvoted 4 times

🗨️ 👤 **vijaya** 2 years ago

A, but may need to remove public IP of ALB from DNS.

upvoted 2 times

🗨️ 👤 **joanneli77** 1 year, 11 months ago

CloudFront requires the origin to be public, therefore ALB must have public IP. That's the root of the problem in this question.

upvoted 2 times

🗨️ 👤 **beznika** 2 years ago

A seems correct.

upvoted 2 times

A company runs several workloads on AWS. The company identifies five AWS Trusted Advisor service quota metrics to monitor in a specific AWS Region. The company wants to receive email notification each time resource usage exceeds 60% of one of the service quotas.


Which solution will meet these requirements?

- A. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.
- B. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- C. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- D. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.

Suggested Answer: C

Community vote distribution


A (100%)

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

Selected Answer: A

it should be SNS instead of SQS

upvoted 7 times

 **Christina666** Highly Voted 11 months, 1 week ago

Selected Answer: A

Amazon CloudWatch allows you to set up alarms on various metrics, including service quota metrics from AWS Trusted Advisor. In this case, the company wants to monitor five Trusted Advisor service quota metrics in a specific AWS Region and receive email notifications when resource usage exceeds 60% of any of these quotas.

Option D is incorrect because it mentions using the AWS Service Health Dashboard to monitor service quota metrics. While the Service Health Dashboard provides information about the overall health of AWS services, it does not provide specific monitoring capabilities for individual service quota metrics as required in this scenario.


upvoted 6 times

 **Mila28** Most Recent 7 months, 2 weeks ago

Selected Answer: A

Service quota metric = Cloudwatch

upvoted 2 times

 **satamex** 9 months, 3 weeks ago

why not D?

upvoted 1 times

 **ctd983** 12 months ago

Selected Answer: A

A is correct

upvoted 3 times

 **jipark** 10 months, 2 weeks ago

"AWS Service Health Dashboard" seems to health check only.

upvoted 1 times

 **CloudHandsOn** 1 year, 6 months ago

Selected Answer: A

Setting up the Alarms, then sending emails via SNS should do the job

upvoted 5 times

A company needs to implement a managed file system to host Windows file shares for users on premises. Resources in the AWS Cloud also need access to the data on these file shares. A SysOps administrator needs to present the user file shares on premises and make the user file shares available on AWS with minimum latency.


What should the SysOps administrator do to meet these requirements?

- A. Set up an Amazon S3 File Gateway.
- B. Set up an AWS Direct Connect connection.
- C. Use AWS DataSync to automate data transfers between the existing file servers and AWS.
- D. Set up an Amazon FSx File Gateway.

Suggested Answer: D

Community vote distribution

D (100%)

 **awsguru1998** Highly Voted 1 year, 4 months ago

Vote for D

To meet the requirements of presenting user file shares on premises and making the file shares available on AWS with minimum latency, the SysOps administrator should set up an Amazon FSx File Gateway. Amazon FSx for Windows File Server provides a fully managed native Microsoft Windows file system so that users can access their files over the Server Message Block (SMB) protocol. The Amazon FSx File Gateway allows access to the file shares from on-premises Windows environments or from Windows instances running in the AWS Cloud with low latency. This solution avoids the need for using AWS Direct Connect, which is more suitable for high-throughput workloads, and eliminates the need for using AWS DataSync to move data between the on-premises file servers and AWS, as the file shares are natively hosted on AWS. Amazon S3 File Gateway, on the other hand, is used for providing a file interface to S3 object storage and does not natively support the SMB protocol, making it unsuitable for this scenario.

upvoted 12 times

 **jipark** Highly Voted 10 months, 2 weeks ago

Selected Answer: D

"Windows file shares for users on premises" = AWS FSx file gateway


upvoted 5 times

 **skiwili** Most Recent 1 year, 6 months ago

Selected Answer: D

Ddddddd

upvoted 4 times

 **beznika** 1 year, 6 months ago

D should do.

upvoted 2 times

A company is hosting applications on Amazon EC2 instances. The company is hosting a database on an Amazon RDS for PostgreSQL DB instance. The company requires all connections to the DB instance to be encrypted.

What should a SysOps administrator do to meet this requirement?

- A. Allow SSL connections to the database by using an inbound security group rule.
- B. Encrypt the database by using an AWS Key Management Service (AWS KMS) encryption key.
- C. Enforce SSL connections to the database by using a custom parameter group.
- D. Patch the database with SSL/TLS by using a custom PostgreSQL extension.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Christina666** Highly Voted 👍 11 months, 1 week ago

Selected Answer: C

C. Enforce SSL connections to the database by using a custom parameter group.

Explanation:

Enforcing SSL connections to the database is the appropriate way to ensure that all connections between the application and the Amazon RDS for PostgreSQL DB instance are encrypted. This can be achieved using a custom parameter group, which allows you to configure specific database settings.

When you enforce SSL connections, the PostgreSQL server will require clients to use SSL when connecting. This ensures that all data transmitted between the application and the database is encrypted, providing an additional layer of security for sensitive information.

Options A, B, and D are not directly related to enforcing SSL connections for the database:

upvoted 5 times

🗳️ 👤 **Christina666** 11 months, 1 week ago

A. Allowing SSL connections to the database through an inbound security group rule would only control network access to the database. While it's a good practice to restrict access, this option alone does not enforce encryption on the connections.

B. Encrypting the database using AWS Key Management Service (AWS KMS) encryption key is a good practice for data-at-rest encryption, but it does not specifically enforce SSL connections for network communication between the application and the database.

D. Patching the database with SSL/TLS by using a custom PostgreSQL extension is not the correct approach. SSL/TLS support is usually built into the PostgreSQL database, and you should not need to patch it with a custom extension for this purpose. Instead, you can enforce SSL connections through the custom parameter group.

upvoted 6 times

🗳️ 👤 **guau** Most Recent 🕒 12 months ago

C, not sure if restart is required but is c

upvoted 2 times

🗳️ 👤 **zolthar_z** 1 year, 6 months ago

Selected Answer: C

yeap, C is the answer: https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithParamGroups.html

upvoted 3 times

🗳️ 👤 **beznika** 1 year, 6 months ago

C looks ok <https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/PostgreSQL.Concepts.General.SSL.html>

upvoted 3 times

A company recently purchased Savings Plans. The company wants to receive email notification when the company's utilization drops below 90% for a given day.

Which solution will meet this requirement?

- A. Create an Amazon CloudWatch alarm to monitor the Savings Plan check in AWS Trusted Advisor. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification when the utilization drops below 90% for a given day.
- B. Create an Amazon CloudWatch alarm to monitor the SavingsPlansUtilization metric under the AWS/SavingsPlans namespace in CloudWatch. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification when the utilization drops below 90% for a given day.
- C. Create a Savings Plans alert to monitor the daily utilization of the Savings Plans. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification when the utilization drops below 90% for a given day.
- D. Use AWS Budgets to create a Savings Plans budget to track the daily utilization of the Savings Plans. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification when the utilization drops below 90% for a given day.

Suggested Answer: C

Community vote distribution

D (83%)

C (17%)

 **Arnaud92** Highly Voted 2 years, 6 months ago

Selected Answer: D

D

<https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-sns-policy.html>

upvoted 8 times

 **hoez** 2 years, 6 months ago

Answer is D Savings Plan Budget notification

<https://www.awsinformation.com/2022/09/26/set-up-notifications-for-low-savings-plan-utilization/>

upvoted 5 times

 **Domdom120** Highly Voted 2 years, 3 months ago


Selected Answer: D

It's D.

You need to use AWS Budgets to manage savings plans. Thresholds and recipients are specified in the settings of your created budget.

<https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>

upvoted 6 times

 **6022a06** Most Recent 10 months, 2 weeks ago

Selected Answer: D

D

<https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>

upvoted 1 times

 **Learning4life** 1 year, 4 months ago

Selected Answer: D

<https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>

upvoted 2 times

 **TareDHakim** 1 year, 5 months ago

D - Saving Plan Budget with SNS

<https://docs.aws.amazon.com/cost-management/latest/userguide/create-savingsplans-budget.html>

upvoted 2 times

 **callspace** 1 year, 8 months ago



Selected Answer: D

From the Savings Plan Guide:

Using budgets

You can use AWS Budgets to set budgets for your Savings Plan utilization, coverage, and costs. You can track your costs as you continue to optimize through AWS.

upvoted 3 times

  **ctd983** 1 year, 12 months ago

Selected Answer: D

D is correct

upvoted 1 times

  **noahsark** 2 years, 2 months ago

Selected Answer: D

Use AWS Budgets to create a Savings Plans budget to track the daily utilization of the Savings Plans. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification when the utilization drops below 90% for a given day.

<https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>

upvoted 2 times

  **defmania00** 2 years, 4 months ago

<https://docs.aws.amazon.com/cost-management/latest/userguide/create-savingsplans-budget.html>

upvoted 3 times

  **awsguru1998** 2 years, 4 months ago

I vote C

Note that D may track the daily utilization of the Savings Plans, but it does not specify how to monitor the utilization and send email notifications when it drops below 90%.

upvoted 2 times

  **skywalker** 2 years, 5 months ago

Selected Answer: D

D Savings Plans budget

upvoted 4 times


  **skywalker** 2 years, 5 months ago

Selected Answer: C

Managing your Savings Plans alerts



You can track your Savings Plans expirations and upcoming queued Savings Plans in Cost Explorer. You can use Savings Plans alerts to receive advance email alerts 1, 7, 30, or 60 days before your Savings Plan expiration date, or in when a commitment is queued for purchase. These notifications also alert you on the expiration date, and can be sent to up to 10 email recipients.

upvoted 1 times

  **Kipalom** 1 year, 6 months ago

Yes, it only notifies you about expiration, but you can't be notified with this feature when the utilization drops below a percentage. For that you need to use AWS Budgets.

upvoted 1 times



  **yeacuz** 2 years, 5 months ago

Selected Answer: C

The answer is definitely D, not C. Here is the documentation on how to create a daily Savings Plan budget with SNS notification:

<https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>

upvoted 3 times

  **yeacuz** 2 years, 5 months ago

The answer is definitely D, not C. Here is the documentation on how to create a daily Savings Plan budget with SNS notification:

<https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>



upvoted 4 times

  **michaldavid** 2 years, 6 months ago

Selected Answer: C

ccccccc

upvoted 1 times

  **yeacuz** 2 years, 5 months ago

The answer is definitely D, not C. Here is the documentation on how to create a daily Savings Plan budget with SNS notification:

<https://docs.aws.amazon.com/savingsplans/latest/userguide/sp-usingBudgets.html>

upvoted 4 times

A company uses an Amazon Simple Queue Service (Amazon SQS) standard queue with its application. The application sends messages to the queue with unique message bodies. The company decides to switch to an SQS FIFO queue.


What must the company do to migrate to an SQS FIFO queue?

- A. Create a new SQS FIFO queue. Turn on content-based deduplication on the new FIFO queue. Update the application to include a message group ID in the messages.
- B. Create a new SQS FIFO queue. Update the application to include the DelaySeconds parameter in the messages.
- C. Modify the queue type from SQS standard to SQS FIFO. Turn off content-based deduplication on the queue. Update the application to include a message group ID in the messages.
- D. Modify the queue type from SQS standard to SQS FIFO. Update the application to send messages with identical message bodies and to include the DelaySeconds parameter in the messages.

Suggested Answer: B

Community vote distribution

A (100%)

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

I think it's A. <https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues-moving.html>
upvoted 9 times

 **callspace** Most Recent 8 months, 4 weeks ago

Selected Answer: A

A it is.

First of all - moving from an Amazon SQS standard queue to a FIFO queue ... You can't convert an existing standard queue into a FIFO queue.
Second - content-based deduplication for the queue (each of your messages has a unique body).
upvoted 4 times

 **jipark** 10 months, 2 weeks ago

Selected Answer: A

why not C-D : Queue type cannot be modified

why not B : DelaySeconds parameter has no relationship with FIFO.

upvoted 3 times

 **Christina666** 11 months, 1 week ago

Selected Answer: A

A.

Create a new SQS FIFO queue: SQS FIFO queues have a different structure and ordering mechanism compared to SQS standard queues, so a new queue needs to be created with the FIFO configuration.

Turn on content-based deduplication on the new FIFO queue: SQS FIFO queues support content-based deduplication. Enabling this feature ensures that identical messages (messages with the same message body) are not duplicated in the queue, which is an essential feature of FIFO queues.
upvoted 3 times

 **Christina666** 11 months, 1 week ago

Update the application to include a message group ID in the messages: In FIFO queues, messages are ordered based on message groups. Each message should include a message group ID, and messages within the same group will be processed in order. If the application is already sending unique message bodies, it might need to be updated to include a message group ID to maintain ordering within the queue.

Option B is incorrect because the DelaySeconds parameter is not relevant to the migration process from standard to FIFO queues. DelaySeconds is a parameter that can be used in both standard and FIFO queues, but it does not impact the migration process.

Option C and Option D are incorrect because modifying the queue type from standard to FIFO or updating the application to send identical message bodies and use DelaySeconds parameter are not appropriate steps for migrating to an SQS FIFO queue.
upvoted 2 times

🗨️ 👤 **csG13** 1 year, 3 months ago

Selected Answer: A

Correct answer is A.

Can't convert an existing queue to a fifo queue, so no C or D. Also FIFO queues only support per-queue delays and not per-message. so it can't be B either.

upvoted 3 times

🗨️ 👤 **braveheart22** 1 year, 3 months ago

A is the correct answer

Every message sent to a FIFO queue requires a message group ID. If you don't need multiple ordered message groups, specify the same message group ID for all your messages.

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues-moving.html>

upvoted 3 times

🗨️ 👤 **nlw** 1 year, 5 months ago

Selected Answer: A

FIFO queues don't support per-message delays, only per-queue delays. If your application sets the same value of the DelaySeconds parameter on each message, you must modify your application to remove the per-message delay and set DelaySeconds on the entire queue instead.

upvoted 3 times

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

Selected Answer: A

I think A as well

upvoted 4 times

A company's SysOps administrator must ensure that all Amazon EC2 Windows instances that are launched in an AWS account have a third-party agent installed. The third-party agent has an .msi package. The company uses AWS Systems Manager for patching, and the Windows instances are tagged appropriately. The third-party agent requires periodic updates as new versions are released. The SysOps administrator must deploy these updates automatically.

Which combination of steps will meet these requirements with the LEAST operational effort? (Choose two.)


- A. Create a Systems Manager Distributor package for the third-party agent.
- B. Make sure that Systems Manager Inventory is configured. If Systems Manager Inventory is not configured, set up a new inventory for instances that is based on the appropriate tag value for Windows.
- C. Create a Systems Manager State Manager association to run the AWS-RunRemoteScript document. Populate the details of the third-party agent package. Specify instance tags based on the appropriate tag value for Windows with a schedule of 1 day.
- D. Create a Systems Manager State Manager association to run the AWS-ConfigureAWSPackage document. Populate the details of the third-party agent package. Specify instance tags based on the appropriate tag value for Windows with a schedule of 1 day.
- E. Create a Systems Manager OpsItem with the tag value for Windows. Attach the Systems Manager Distributor package to the OpsItem. Create a maintenance window that is specific to the package deployment. Configure the maintenance window to cover 24 hours a day.

Suggested Answer: AD

Community vote distribution

AD (65%)

AC (35%)

 **zolthar_z** Highly Voted 2 years, 6 months ago

Selected Answer: AD

A and D <https://docs.aws.amazon.com/systems-manager/latest/userguide/distributor-working-with-packages-deploy.html>
upvoted 5 times

 **xSohox** Highly Voted 1 year, 10 months ago

Selected Answer: AD

Correct answer is AD:

You can just read a description of each SSM document:

AWS-ConfigureAWSPackage

"Install or uninstall a Distributor package. You can install the latest version, default version, or a version of the package you specify."

AWS-RunRemoteScript

"Execute scripts stored in a REMOTE location. The following remote locations are currently supported: GitHub (public and private) and Amazon S3 (S3). The following script types are currently supported: #! support on Linux and file associations on Windows."

So we need to create a Distributor package and then create State Manager association to run the AWS-ConfigureAWSPackage
upvoted 5 times

 **walala97** 1 year, 8 months ago

"AWS-RunRemoteScript "for collaboration with third-party agents..The "AWS ConfigureAWSPackage" is mainly used to configure the AWS Management Agent (SSM Agent) to ensure that communication between EC2 instances and AWS Systems Manager is working properly. This script is typically used to ensure that AWS's own agent (i.e. SSM Agent) is functioning properly.but the option A said we need the third-party ,so the answer is AC

upvoted 2 times

 **Debugs_Bunny** 1 year, 6 months ago

key is "LEAST operational effort". Answer is AD

upvoted 1 times

 **numark** Most Recent 7 months ago

Selected Answer: AD

With State Manager, a capability of AWS Systems Manager, you can define and maintain consistent configuration of your operating system and software. Using the AWS-ConfigureAWSPackage document in a State Manager association allows you to install, update, or remove packages. By

specifying instance tags based on the appropriate tag value for Windows with a schedule (e.g., every 1 day), you can automate the deployment of updates to the agent across all your managed Windows EC2 instances. The AWS-RunRemoteScript document can be used to download and execute scripts on EC2 instances, but it requires the management and storage of the script separately and would involve more operational effort in handling package updates.

upvoted 1 times

🗨️ 👤 **March2023** 1 year, 3 months ago

Selected Answer: AD

a & d instead of C

upvoted 3 times

🗨️ 👤 **Hatem08** 1 year, 6 months ago

Selected Answer: AC

AC makes more sense for me as it is third party package

upvoted 1 times

🗨️ 👤 **telosd** 1 year, 5 months ago

AD:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/distributor-working-with-packages-deploy.html#distributor-deploy-sm-pkg-console>

upvoted 3 times

🗨️ 👤 **jipark** 1 year, 10 months ago

Selected Answer: AC

why not D : AWS-ConfigureAWSPackage is for AWS installation.

why C : AWS-RunRemoteScript is for 3rd party,

upvoted 3 times

🗨️ 👤 **Christina666** 1 year, 11 months ago

Selected Answer: AD

Step A: Create a Systems Manager Distributor package for the third-party agent.

Systems Manager Distributor allows you to package and distribute software and files to your instances using SSM. By creating a Distributor package for the third-party agent's .msi package, you can centrally manage its installation and updates across EC2 instances.

Step D: Create a Systems Manager State Manager association to run the AWS-ConfigureAWSPackage document. Populate the details of the third-party agent package. Specify instance tags based on the appropriate tag value for Windows with a schedule of 1 day.

AWS-ConfigureAWSPackage document is used to install software packages on EC2 instances. By creating a State Manager association with AWS-ConfigureAWSPackage document and specifying the instance tags based on the appropriate tag value for Windows, you can ensure that the third-party agent package is deployed on the instances automatically.

upvoted 2 times

🗨️ 👤 **Christina666** 1 year, 11 months ago

Options B and C are not necessary for this scenario as they are not the most efficient or direct ways to achieve the goal. Option E is also not the best approach, as using an OpsItem and creating a maintenance window for a 24-hour deployment is unnecessary complexity for this requirement.

upvoted 2 times

🗨️ 👤 **guau** 1 year, 12 months ago

Selected Answer: AC

AC third party package!

upvoted 3 times

🗨️ 👤 **Gomer** 2 years, 2 months ago

Selected Answer: AD

Isn't this example the answer to the question of "C" over "D"? 1st step is to create an SSM formatted "package" using any install files (AWS, MS, or whoever). Then you use the AWS-ConfigureAWSPackage "document" to create the association.

Create Association to update a Distributor package on a schedule without taking application offline

aws ssm create-association --name "AWS-ConfigureAWSPackage" [...]

<https://docs.aws.amazon.com/systems-manager/latest/userguide/distributor-working-with-packages-deploy.html#distributor-smupdate-pkg-cli>

upvoted 3 times

🗨️ 👤 **noahsark** 2 years, 2 months ago

Selected Answer: AD

Create a Systems Manager Distributor package for the third-party agent.

Create a Systems Manager State Manager association to run the AWS-ConfigureAWSPackage document. Populate the details of the third-party agent package. Specify instance tags based on the appropriate tag value for Windows with a schedule of 1 day.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/distributor-working-with-packages-create.html>

upvoted 2 times

🗨️ 👤 **JamesF92** 1 year, 10 months ago

this link from Noahsark shows *.msi is a supported package type for SSM. We have a .msi.

So D is correct (not C). AWS-RunRemoteScript (C) is for grabbing scripts from Github, etc.

upvoted 1 times

🗨️ 👤 **vhernan** 2 years, 2 months ago

Selected Answer: AC

ad is correct

upvoted 1 times

🗨️ 👤 **a1971h** 2 years, 2 months ago

Selected Answer: AC

AC is more correct

upvoted 1 times

🗨️ 👤 **Vivec** 2 years, 3 months ago

Selected Answer: AC

Option D is incorrect because AWS-ConfigureAWSPackage document is used to install and configure AWS packages, not third-party packages. It is used to install packages from the Amazon Linux Extras library, AWS services such as the SSM agent, and third-party software provided by AWS Partners.

To install the third-party agent on EC2 instances, the best approach is to create a Systems Manager Distributor package for the third-party agent, as mentioned in option A. The Distributor package allows you to centrally manage software deployment across your Amazon EC2 instances or on-premises servers. You can use this feature to deploy both AWS and third-party software. Once you create the package, you can use State Manager associations to schedule and deploy the package to EC2 instances based on instance tags, as mentioned in option C.

upvoted 4 times

🗨️ 👤 **vn_thanhtung** 2 years, 3 months ago

I think ans is A and D because "requirements with the LEAST operational effort".

upvoted 1 times

🗨️ 👤 **vn_thanhtung** 2 years, 3 months ago

sorry for this confusion I was too hasty to give an answer when I didn't read the question carefully the answer should be AC.

upvoted 1 times

🗨️ 👤 **Domdom120** 2 years, 4 months ago

Selected Answer: AD

A,D

Correct reference provided by zolthar. I have no idea what awsguru is referencing that says AWS-ConfigureAWSPackage can only be used for AWS-provided packages, as that would require a tremendous overhead on AWS for providing packages on an insane amount of available software. This is very similar to SCCM in which you build your package yourself with the software you want and set the install targets/requirements.

upvoted 2 times

🗨️ 👤 **awsguru1998** 2 years, 4 months ago

AC

Not D, it is not the AWS-ConfigureAWSPackage document. The AWS-ConfigureAWSPackage document is used to install, uninstall, or update an AWS-provided package or application. In this scenario, the SysOps administrator is trying to install a third-party agent, which is not an AWS-provided package or application. Therefore, the AWS-RunRemoteScript document is the appropriate choice to install and update the third-party agent.

upvoted 3 times

🗨️ 👤 **skiwili** 2 years, 6 months ago

Selected Answer: AD

AD for sure

upvoted 2 times

A company runs hundreds of Amazon EC2 instances in a single AWS Region. Each EC2 instance has two attached 1 GiB General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volumes. A critical workload is using all the available IOPS capacity on the EBS volumes.

According to company policy, the company cannot change instance types or EBS volume types without completing lengthy acceptance tests to validate that the company's applications will function properly. A SysOps administrator needs to increase the I/O performance of the EBS volumes as quickly as possible.

Which action should the SysOps administrator take to meet these requirements?


- A. Increase the size of the 1 GiB EBS volumes.
- B. Add two additional elastic network interfaces on each EC2 instance.
- C. Turn on Transfer Acceleration on the EBS volumes in the Region.
- D. Add all the EC2 instances to a cluster placement group.

Suggested Answer: A

Community vote distribution

A (84%)

Other

 **Domdom120** Highly Voted 1 year, 10 months ago

Selected Answer: A

A. 100% positive.


"With Amazon EBS Elastic Volumes, you can increase the volume size, change the volume type, or adjust the performance of your EBS volumes. If your instance supports Elastic Volumes, you can do so without detaching the volume or restarting the instance. This enables you to continue using your application while the changes take effect."

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modify-volume.html>

They have Elastic Volumes in place (per the question) and that's exactly why it is specified in the question. As others have mentioned, increasing the volume size increases IOPS, up to the volume type max. For gp2, you can have a volume size of 1 GiB - 16 TiB with a max IOPS of 16,000 for the 16 TiB volume size.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html>

upvoted 8 times

 **jipark** 1 year, 4 months ago

I don't know why, but

"By increasing the size of the EBS volumes, you can effectively increase the IOPS capacity available to your instances. "

upvoted 5 times


 **johnson_chao** Most Recent 9 months ago

Selected Answer: A

Answer is A

ref : <https://docs.aws.amazon.com/ebs/latest/userguide/general-purpose.html>


upvoted 2 times

 **Gomer** 1 year, 8 months ago

Selected Answer: A

I think A is the correct answer based on the question. However, in the real world, I'd be dubious that changing the volume size (and file system size) would increase IOPS in this case, at least not instantly and maybe not significantly until you actually start writing to the new blocks in a big way. I'd be willing to bet, I/O for existing blocks is going to remain unchained. I'd bet your not going to get instantly greater IOPS just because you resized the volume.



upvoted 1 times

 **Vivec** 1 year, 9 months ago

Selected Answer: A



The IOPS performance of Amazon EBS volumes is directly proportional to the size of the volume. Therefore, increasing the size of the EBS volumes is the most straightforward way to increase the IOPS performance of the volumes.

upvoted 3 times

  **Vivec** 1 year, 9 months ago

Options B, C, and D are not relevant to the situation described in the question. Adding additional network interfaces, turning on Transfer Acceleration on EBS volumes, or adding EC2 instances to a placement group will not directly improve the I/O performance of the EBS volumes.

upvoted 1 times

  **Brokdar** 1 year, 10 months ago

Selected Answer: A

A is correct, since increasing the size of the EBS volume doesn't change the volume type and it will definitely increase the IOPS. B, C and D don't make sense for this question.

upvoted 4 times

  **joanneli77** 1 year, 11 months ago

A - I'm not changing EBS volume TYPE, I'm changing EBS volume size. Cluster placement group does not affect IOPS to a local drive on individual instances, but can optimize network communication latency.

upvoted 1 times

  **wooyourdaddy** 1 year, 11 months ago

Selected Answer: D

D. Add all the EC2 instances to a cluster placement group.

100% not A: Amazon S3 Transfer Acceleration is a bucket-level feature that enables fast, easy, and secure transfers of files over long distances between your client and an S3 bucket. Transfer Acceleration is designed to optimize transfer speeds from across the world into S3 buckets.

Cluster placement groups

A cluster placement group is a logical grouping of instances within a single Availability Zone. A cluster placement group can span peered virtual private networks (VPCs) in the same Region. Instances in the same cluster placement group enjoy a higher per-flow throughput limit for TCP/IP traffic and are placed in the same high-bisection bandwidth segment of the network.

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



upvoted 1 times

  **wooyourdaddy** 1 year, 11 months ago

100% not A, changing the size or adding size would require changing the config, which would require testing.



100 % not C: A: Amazon S3 Transfer Acceleration is a bucket-level feature that enables fast, easy, and secure transfers of files over long distances between your client and an S3 bucket. Transfer Acceleration is designed to optimize transfer speeds from across the world into S3 buckets.

upvoted 1 times

  **Brokdar** 1 year, 10 months ago

Increasing the size does not require changing any config, just a partition resize (on Windows) to take advantage of the new size or similar on LVM (on Linux). This doesn't change the EBS volume type, and it will definitely increase the IOPS. So, 100% A.

upvoted 2 times

  **Fatoch** 1 year, 11 months ago

For me C is correct. Performance should be quick

upvoted 1 times

  **skiwili** 2 years ago

Selected Answer: A

A would be the correct answer

upvoted 3 times

  **zolphar_z** 2 years ago

Answer is A, more capacity = more IOPS

upvoted 3 times

  **michaldavid** 2 years ago

Selected Answer: D

I stand corrected this is D

upvoted 2 times

  **beznika** 2 years ago

I think A is correct. C is for S3 and D doesn't make any sense. Moving instances to placement groups would require acceptance tests.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/requesting-ebs-volume-modifications.html>

upvoted 2 times

  **tt79** 2 years ago

Transfer acceleration is used for uploading data to S3.

D is correct.

upvoted 1 times

  **michaldavid** 2 years ago

Selected Answer: C

This is C to me

upvoted 1 times

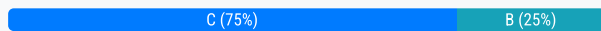
A company needs to deploy a new workload on AWS. The company must encrypt all data at rest and must rotate the encryption keys once each year. The workload uses an Amazon RDS for MySQL Multi-AZ database for data storage.

Which configuration approach will meet these requirements?

- A. Enable Transparent Data Encryption (TDE) in the MySQL configuration file. Manually rotate the key every 12 months.
- B. Enable RDS encryption on the database at creation time by using the AWS managed key for Amazon RDS.
- C. Create a new AWS Key Management Service (AWS KMS) customer managed key. Enable automatic key rotation. Enable RDS encryption on the database at creation time by using the KMS key.
- D. Create a new AWS Key Management Service (AWS KMS) customer managed key. Enable automatic key rotation. Enable encryption on the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the RDS DB instance.

Suggested Answer: C

Community vote distribution



🗳️ 👤 **Rado_Piatek** 1 month ago

Selected Answer: C

Both B and C are valid option, if the question was to choose the least effort option i would go with B, but C seems more rational in this case.

AWS managed keys rotate automatically so you don't to enable rotation additionally, in fact you cannot change it to NOT rotate.

upvoted 1 times

🗳️ 👤 **TareDHakim** 12 months ago

Selected Answer: B

B is easiest!

Why would you choose C, a custom key with management overhead when you can just have AWS encrypt and rotate the key using default Encryption settings?

upvoted 2 times

🗳️ 👤 **confusedyeti69** 10 months ago

B doesn't say "Enable automatic key rotation".

upvoted 5 times

🗳️ 👤 **Christina666** 1 year, 5 months ago

Selected Answer: C

Automatic key rotation is supported only on symmetric encryption KMS keys. You cannot enable automatic rotation of asymmetric KMS keys, HMAC KMS keys, KMS keys with imported key material, or KMS keys in a custom key store. To enable or disable automatic rotation of a set of related multi-Region keys, set the property on the primary key.

upvoted 4 times

🗳️ 👤 **CodePoet** 1 year, 12 months ago

Selected Answer: C

Obviously C

upvoted 2 times

🗳️ 👤 **michaldavid** 2 years ago

Selected Answer: C

Agree with C

upvoted 2 times

A company has an application that is deployed to two AWS Regions in an active-passive configuration. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The instances are in an Amazon EC2 Auto Scaling group in each Region. The application uses an Amazon Route 53 hosted zone for DNS. A SysOps administrator needs to configure automatic failover to the secondary Region.

What should the SysOps administrator do to meet these requirements?

- A. Configure Route 53 alias records that point to each ALB. Choose a failover routing policy. Set Evaluate Target Health to Yes.
- B. Configure CNAME records that point to each ALB. Choose a failover routing policy. Set Evaluate Target Health to Yes.
- C. Configure Elastic Load Balancing (ELB) health checks for the Auto Scaling group. Add a target group to the ALB in the primary Region. Include the EC2 instances in the secondary Region as targets.
- D. Configure EC2 health checks for the Auto Scaling group. Add a target group to the ALB in the primary Region. Include the EC2 instances in the secondary Region as targets.

Suggested Answer: A

Community vote distribution


A (100%)

 **Aamee** 7 months, 4 weeks ago

Selected Answer: A

Hint: Alias record gets associated with ALB so def. option A!..

upvoted 1 times

 **10cc6ba** 11 months, 2 weeks ago

C

C. Store the data in S3 Standard for the first 90 days. Set up an S3 Lifecycle rule to move the data to S3 Glacier Flexible Retrieval after 90 days.

Explanation:

First 90 Days:

S3 Standard provides high availability and low latency, which meets the requirement of providing access to the data in milliseconds during the first 90 days when the data is infrequently accessed but must remain highly available.

After 90 Days:

S3 Glacier Flexible Retrieval is a cost-effective storage class designed for long-term data archiving with retrieval times typically within a few hours, meeting the requirement of retrieval time under 5 hours. This option provides a significant cost saving compared to S3 Standard while still ensuring that data can be retrieved relatively quickly.

upvoted 1 times

 **Christina666** 1 year, 11 months ago

Selected Answer: A

A Route 53 failover routing

upvoted 4 times

 **marco25** 2 years, 3 months ago

Selected Answer: A

Alias allow duplicate record name which is the case here

upvoted 3 times

 **Vivec** 2 years, 3 months ago

Selected Answer: A

To configure automatic failover to the secondary Region for an application that is deployed to two AWS Regions in an active-passive configuration, the following steps should be taken:

Configure Route 53 alias records that point to each ALB in the two Regions.

Choose a failover routing policy, such as Failover or Geolocation.

Set Evaluate Target Health to Yes to ensure that Route 53 only responds to DNS queries with healthy ALB endpoints.



upvoted 3 times

  **joanneli77** 2 years, 4 months ago

Selected Answer: A

Active / Standby = one region active at a time. Health checks and Route53.

upvoted 3 times

  **Gomer** 2 years, 2 months ago

I believe it's Active/Active unless app in one region is inaccessible for whatever reason. If that happens, Route 53 directs all user traffic to the running region. Yes, access may be slower for some users, but slower is better than dead in the water.

upvoted 1 times

A company is implementing a monitoring solution that is based on machine learning. The monitoring solution consumes Amazon EventBridge (Amazon CloudWatch Events) events that are generated by Amazon EC2 Auto Scaling. The monitoring solution provides detection of anomalous behavior such as unanticipated scaling events and is configured as an EventBridge (CloudWatch Events) API destination.

During initial testing, the company discovers that the monitoring solution is not receiving events. However, Amazon CloudWatch is showing that the EventBridge (CloudWatch Events) rule is being invoked. A SysOps administrator must implement a solution to retrieve client error details to help resolve this issue.


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

- A. Create an EventBridge (CloudWatch Events) archive for the event pattern to replay the events. Increase the logging on the monitoring solution. Use replay to invoke the monitoring solution. Examine the error details.
- B. Add an Amazon Simple Queue Service (Amazon SQS) standard queue as a dead-letter queue for the target. Process the messages in the dead-letter queue to retrieve error details.
- C. Create a second EventBridge (CloudWatch Events) rule for the same event pattern to target an AWS Lambda function. Configure the Lambda function to invoke the monitoring solution and to record the results to Amazon CloudWatch Logs. Examine the errors in the logs.
- D. Configure the EventBridge (CloudWatch Events) rule to send error messages to an Amazon Simple Notification Service (Amazon SNS) topic.

Suggested Answer: B

Community vote distribution

A (53%) B (48%)

 **Domdom120** Highly Voted 2 years, 4 months ago

Selected Answer: B

B.

The questions states: "A SysOps administrator must implement a solution to retrieve client error details to help resolve this issue."


Supporting answer:

"Amazon SQS supports dead-letter queues (DLQ), which other queues (source queues) can target for messages that can't be processed (consumed) successfully. Dead-letter queues are useful for debugging your application or messaging system because they let you isolate unconsumed messages to determine why their processing doesn't succeed."

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-dead-letter-queues.html>

I'm really disappointed in the comments sections for this exam in particular. If you're going to take the time to comment, please include reference documentation that supports your choice.

upvoted 17 times

 **robotgeek** 1 year, 8 months ago

Agree with providing sources but disagree that the dead-letter would allow to "retrieve client error details", the SQS deadletter would hold the message, not any information about the error. This is exactly what option A does, allow you to automatically replay messages and see the error logging in the external (to AWS) application


upvoted 2 times

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

Selected Answer: A

Answer is absolutely "A" based on the links already provided. You need to replay the events, and try and identify why the monitoring app isn't able to contact EventBridge API. If you replace the monitoring app with SQS que, your just going to log Auto Scaling events. That isn't going to help you figure out why monitoring app isn't getting these same events through the EventBridge API. The SQS "B" answer may be easier than "A", but it isn't going to do any good on debugging EventBridge API access issue (probably a permissions/role/policy issue)

upvoted 8 times

 **jipark** 1 year, 10 months ago

I agree, DLQ shows just Queue connection error.

upvoted 1 times

  **hinda** Most Recent 2 months ago

Selected Answer: B

EventBridge supports DLQs for targets like API destinations.

DLQs capture events that fail delivery due to issues like authentication errors, bad responses (4xx/5xx), etc.


upvoted 1 times

  **numark** 7 months ago

Selected Answer: B

B>>This option is most suitable because Amazon EventBridge allows you to configure a dead-letter queue (DLQ) to handle events that fail to be delivered to the target. By directing these failed events to an Amazon SQS queue, you can store the events and analyze them to determine why the delivery failed, which may include client error details. This solution requires minimal operational effort as it leverages AWS's built-in mechanisms for error handling without the need to create additional resources or logging mechanisms. Creating an archive and replaying events is a more complex solution that might be necessary if you need to resend the events after fixing the issue. However, it doesn't directly address the requirement to retrieve error details and involves more operational steps compared to setting up a dead-letter queue.



upvoted 2 times

  **rex3** 9 months ago

Selected Answer: B

Dead-letter queues (DLQ) are a feature in EventBridge that allows failed events (such as client errors when invoking an API destination) to be sent to a dead-letter queue (in this case, an Amazon SQS queue).

upvoted 1 times

  **VerRi** 11 months, 1 week ago

Selected Answer: A

EventBridge Archive can retain event copies and replay them for debugging.

upvoted 2 times

  **AWSdeveloper08** 1 year, 10 months ago

Selected Answer: A

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-archive.html>

upvoted 6 times

  **Christina666** 1 year, 11 months ago



Selected Answer: A

EventBridge can now archive and replay events:

You can now create an encrypted archive of the events published to an event bus. You can archive all events, or filter them using the same pattern matching syntax used by EventBridge rules. You can store events indefinitely, or set up a retention period after which older events are automatically removed from the archive.

You can also replay the events stored in an archive. Events are replayed to all rules defined for the event bus (but not to managed rules created by other AWS services) or to the rules you specify. Replayed events contain an extra replay-name field in case you need to recognize them. When starting a replay, you define a time frame, and only events within that time frame are replayed. Currently, you can only replay events to the same event bus from which they were archived.

upvoted 4 times

  **RayHK** 1 year, 11 months ago

why not C? <https://docs.aws.amazon.com/lambda/latest/dg/services-cloudwatchevents.html>

upvoted 1 times

  **Vivec** 2 years, 2 months ago

Selected Answer: B

Based on the requirements to retrieve client error details with the least operational effort, option B would be the best solution. Adding an SQS standard queue as a dead-letter queue for the target will capture and store any failed messages, including error details. The messages can then be processed later to retrieve the error details, without requiring any additional configurations or modifications to the monitoring solution.

Option A requires creating an EventBridge (CloudWatch Events) archive and replaying events, increasing logging, and then examining the error details. This solution requires additional effort and may not necessarily capture the exact error details that occurred.

upvoted 2 times



  **thatTeller30** 2 years, 3 months ago

Selected Answer: A

Im in my opinion its A - <https://aws.amazon.com/blogs/aws/new-archive-and-replay-events-with-amazon-eventbridge/>

mainly because in the question it states they use the EventBridge API, meaning nothing goes into SQS and its unrelated to the question. By Default there's not invocation between EventBridge and SQS. U can however, create an Invoke Rule from EventBridge into SQS.

upvoted 2 times

  **braveheart22** 2 years, 3 months ago

A is the correct answer from my point of view.

After carefully reading the question, I agree with "yeacuz" argument from the link below.

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-archive.html>

In EventBridge, you can create an archive of events so that you can easily replay them at a later time. For example, you might want to replay events to recover from errors or to validate new functionality in your application.

upvoted 1 times

  **awsguru1998** 2 years, 4 months ago



Option D, which is to configure the EventBridge (CloudWatch Events) rule to send error messages to an Amazon Simple Notification Service (Amazon SNS) topic, would be more efficient in terms of operational effort compared to Option C. This is because Option D requires only a simple configuration of the EventBridge (CloudWatch Events) rule and does not require the creation of a second EventBridge rule or a Lambda function. The error messages can be easily retrieved from the Amazon SNS topic and examined. I apologize for the error in my previous answer.

upvoted 1 times

  **defmania00** 2 years, 4 months ago

Why would you sent error messages to SNS? Doesn't make any sense to me. An Amazon SNS topic is a logical access point that acts as a communication channel. A topic lets you group multiple endpoints (such as AWS Lambda, Amazon SQS, HTTP/S, or an email address).

upvoted 1 times

  **yeacuz** 2 years, 5 months ago

Selected Answer: A

Answer is A:

"In EventBridge, you can create an archive of events so that you can easily replay them at a later time. For example, you might want to replay events to recover from errors or to validate new functionality in your application."

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-archive.html>

upvoted 4 times

  **zolthar_z** 2 years, 5 months ago

Selected Answer: B

Answer is B

upvoted 3 times

A company is storing backups in an Amazon S3 bucket. The backups must not be deleted for at least 3 months after the backups are created.

What should a SysOps administrator do to meet this requirement?

- A. Configure an IAM policy that denies the s3:DeleteObject action for all users. Three months after an object is written, remove the policy.
- B. Enable S3 Object Lock on a new S3 bucket in compliance mode. Place all backups in the new S3 bucket with a retention period of 3 months.
- C. Enable S3 Versioning on the existing S3 bucket. Configure S3 Lifecycle rules to protect the backups.
- D. Enable S3 Object Lock on a new S3 bucket in governance mode. Place all backups in the new S3 bucket with a retention period of 3 months.

Suggested Answer: B

Community vote distribution

B (100%)

 **Domdom120** Highly Voted 10 months ago

Selected Answer: B

B. Compliance mode is required for this situation. Comparison and reference below:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html>

In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. You can also use governance mode to test retention-period settings before creating a compliance-mode retention period.

In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. 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 9 times

 **zolthar_z** Highly Voted 12 months ago

Selected Answer: B

Answer is B, governance mode allows some users to delete/modify the data

upvoted 8 times

 **Gomer** Most Recent 8 months, 1 week ago

Selected Answer: B

Nobody on customer side including root can delete objects before they expire (unless maybe if you delete the entire account). I do know there is a "force" option for deleting buckets with existing objects. Would be interesting to see if that would do it. I'm sure Amazon can always delete files if they have to (and customer insisted they do it).

upvoted 2 times

 **awsguru1998** 10 months, 2 weeks ago

B The S3 Object Lock feature in compliance mode can be used to enforce a retention period for objects in the bucket. This ensures that the backups are protected from deletion for at least 3 months after they are created, which meets the requirement. The retention period can be set at the bucket level, or at the object level.

upvoted 2 times

A SysOps administrator needs to track the costs of data transfer between AWS Regions. The SysOps administrator must implement a solution to send alerts to an email distribution list when transfer costs reach 75% of a specific threshold.

What should the SysOps administrator do to meet these requirements?

- A. Create an AWS Cost and Usage Report. Analyze the results in Amazon Athena. Configure an alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when costs reach 75% of the threshold. Subscribe the email distribution list to the topic.
- B. Create an Amazon CloudWatch billing alarm to detect when costs reach 75% of the threshold. Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the email distribution list to the topic.
- C. Use AWS Budgets to create a cost budget for data transfer costs. Set an alert at 75% of the budgeted amount. Configure the budget to send a notification to the email distribution list when costs reach 75% of the threshold.
- D. Set up a VPC flow log. Set up a subscription filter to an AWS Lambda function to analyze data transfer. Configure the Lambda function to send a notification to the email distribution list when costs reach 75% of the threshold.

Suggested Answer: C

Community vote distribution

C (100%)

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

Selected Answer: C

I was very confused by B and C, because they both do alerts. However, only "Budgets" lets yo select specific services and regions, and alarm based on forecast expenses. I made up a table of the differences between CW Billing Alarms and Budgets alarms, which is too big to post. However, these links helped gel my thoughts on this:

<https://aws.amazon.com/blogs/aws/aws-budgets-update-track-cloud-costs-and-usage/>

<https://stackoverflow.com/questions/67680601/whats-the-difference-between-alarm-budget-and-cloudwatch-alarms-billing>

upvoted 6 times

 **jipark** Most Recent 10 months, 2 weeks ago

Selected Answer: C

as to cost, let's use "AWS Budget" on this exam.

upvoted 4 times

 **Christina666** 11 months, 1 week ago

Selected Answer: C

Option A (Create an AWS Cost and Usage Report. Analyze the results in Amazon Athena. Configure an alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when costs reach 75% of the threshold. Subscribe the email distribution list to the topic.) is a bit complex and not directly related to tracking data transfer costs between AWS Regions. AWS Cost and Usage Reports and Amazon Athena are typically used for analyzing detailed cost and usage data, but they are not the most straightforward solution for this specific requirement.

Option B (Create an Amazon CloudWatch billing alarm to detect when costs reach 75% of the threshold. Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the email distribution list to the topic.) is focused on billing alarms, which can track costs, but it may not specifically address data transfer costs between AWS Regions.

upvoted 4 times

 **Christina666** 11 months, 1 week ago

Option D (Set up a VPC flow log. Set up a subscription filter to an AWS Lambda function to analyze data transfer. Configure the Lambda function to send a notification to the email distribution list when costs reach 75% of the threshold.) is an overcomplicated approach. VPC flow logs are primarily used for monitoring network traffic, and while they can be used for analyzing data transfer, they are not the most suitable option for tracking costs and sending cost-related alerts.

Option C (Use AWS Budgets to create a cost budget for data transfer costs. Set an alert at 75% of the budgeted amount. Configure the budget to send a notification to the email distribution list when costs reach 75% of the threshold.) is the best option. AWS Budgets is a service designed explicitly for tracking and managing costs. By setting up a budget specifically for data transfer costs between AWS Regions and configuring an alert at 75% of the threshold, the SysOps administrator can receive timely alerts through email when the costs approach the defined threshold.

upvoted 3 times

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

Selected Answer: C

maybe C

Use AWS Budgets to create a cost budget for data transfer costs. Set an alert at 75% of the budgeted amount. Configure the budget to send a notification to the email distribution list when costs reach 75% of the threshold.

<https://aws.amazon.com/blogs/aws/aws-budgets-update-track-cloud-costs-and-usage/>

upvoted 2 times

🗨️ 👤 **awsguru1998** 1 year, 4 months ago

B is the most straightforward and easiest option, as it only requires creating an Amazon CloudWatch billing alarm and configuring it to send a message to an Amazon SNS topic when costs reach 75% of the threshold. Option C would achieve the same result, but it involves creating a budget and setting up alerts, which may be more complex to set up and maintain compared to creating a CloudWatch billing alarm.

upvoted 2 times

🗨️ 👤 **dangji** 1 year, 5 months ago

Selected Answer: C

certainly C

upvoted 2 times

🗨️ 👤 **zolthar_z** 1 year, 5 months ago

Selected Answer: C

Ans is C

upvoted 2 times

A company needs to archive all audit logs for 10 years. The company must protect the logs from any future edits.



Which solution will meet these requirements?

- A. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. Configure AWS Key Management Service (AWS KMS) encryption.
- B. Store the data in an Amazon S3 Glacier vault. Configure a vault lock policy for write-once, read-many (WORM) access.
- C. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Configure server-side encryption.
- D. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Configure multi-factor authentication (MFA).

Suggested Answer: B

Community vote distribution

B (100%)

  **aws guru1998** Highly Voted 10 months, 2 weeks ago

B. Store the data in an Amazon S3 Glacier vault. Configure a vault lock policy for write-once, read-many (WORM) access. This will ensure that the data is stored securely with encryption and that it cannot be modified or deleted once it has been stored in the vault. The write-once, read-many (WORM) access provided by a vault lock policy helps ensure data immutability and compliance with regulatory requirements for long-term data retention.

upvoted 6 times



  **dangji** Most Recent 11 months, 4 weeks ago

Selected Answer: B

B

https://docs.aws.amazon.com/zh_tw/AmazonS3/latest/userguide/object-lock.html

upvoted 2 times

  **zolthar_z** 12 months ago

Selected Answer: B

Ans is B

upvoted 2 times