



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

A company has a mobile application that makes HTTP API calls to an Application Load Balancer (ALB). The ALB routes requests to an AWS Lambda function. Many different versions of the application are in use at any given time, including versions that are in testing by a subset of users. The version of the application is defined in the user-agent header that is sent with all requests to the API.

After a series of recent changes to the API, the company has observed issues with the application. The company needs to gather a metric for each API operation by response code for each version of the application that is in use. A DevOps engineer has modified the Lambda function to extract the API operation name, version information from the user-agent header and response code.

Which additional set of actions should the DevOps engineer take to gather the required metrics?


- A. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name. Specify response code and application version as dimensions for the metric.
- B. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group. Configure a CloudWatch Logs Insights query to populate CloudWatch metrics from the log lines. Specify response code and application version as dimensions for the metric.
- C. Configure the ALB access logs to write to an Amazon CloudWatch Logs log group. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadata. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name. Specify response code and application version as dimensions for the metric.
- D. Configure AWS X-Ray integration on the Lambda function. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version number. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatch. Specify response code and application version as dimensions for the metric.

Suggested Answer: B

Community vote distribution

A (93%)

5%

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

Selected Answer: A


Option A: This option is the most efficient way to gather the required metrics. It does not require any additional infrastructure and can be easily implemented.

Option B: This option is more complex than Option A and requires configuring a CloudWatch Logs Insights query. This can be more time-consuming to set up and can be less efficient if the query is not optimized.

Option C: This option requires configuring the ALB access logs to write to CloudWatch Logs. This can add additional latency to the requests.

Option D: This option requires configuring AWS X-Ray integration. This is a more complex solution that is not necessary in this case.

upvoted 5 times

 **festusbetd** Most Recent 2 weeks, 4 days ago

Selected Answer: A

Why Option A is Correct:

It directly addresses the requirement to gather metrics for each API operation by response code and application version.

CloudWatch Logs and metric filters are designed for this purpose and are easy to implement.

It provides a scalable and cost-effective solution.


upvoted 1 times

 **maikerusukofirudo** 3 months, 1 week ago

Selected Answer: A

That's right

upvoted 1 times

 **chan123** 3 months, 2 weeks ago

Selected Answer: D

dsfdsf

upvoted 1 times

 **c3518fc** 5 months, 2 weeks ago

Selected Answer: A

The other options are either incomplete or involve unnecessary complexity:

Option B requires using CloudWatch Logs Insights queries, which may introduce additional complexity and potential performance overhead.

Option C involves modifying the ALB access logs, which may not provide the required level of granularity or flexibility for capturing the application version information.

Option D requires integrating AWS X-Ray, which is primarily designed for distributed tracing and may not be necessary for this specific use case of gathering metrics by response code and application version.

Therefore, option A is the most appropriate and straightforward solution for the given requirements.

upvoted 4 times

  **nothinmuch** 1 year ago

Selected Answer: B

The answer is b based on the scenario.

When to Choose Which

Choose A (Metric Filters) if you need basic metrics with straightforward patterns (e.g., counting occurrences of specific API operations and response codes).

Choose B (Insights Queries) if you require more complex metric calculations, such as:

Aggregations (averages, sums, etc.) over time

Filtering metrics based on conditions within the logs

Creating custom metrics not directly defined in log lines

upvoted 2 times

  **Gowtham5798** 1 year, 1 month ago

The correct answer is A.

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

A is correct: Because only need to gather metric but not parse log or advanced analyzing the log.

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

A definitely

upvoted 1 times

  **hoakhanh281** 1 year, 2 months ago

Cloudwatch log metrics filter can apply on pattern and populate metrics, CW insight significant to get it. B should be correct one



upvoted 1 times

  **d262e67** 1 year, 2 months ago

Selected Answer: A

CloudWatch Logs Insights is a powerful tool for analyzing log data, but it's more suited for ad-hoc exploration and troubleshooting rather than continuous metric collection and monitoring.

upvoted 3 times

  **DucSiu** 1 year, 2 months ago

The company needs to gather a metric for each API operation by response code for each version of the application that is in use


=> A

upvoted 1 times

  **wem** 1 year, 3 months ago

Option A seems to be the most straightforward and effective method. By writing the required information as a log line to CloudWatch Logs and configuring a metric filter to increment metrics based on this data, the company can efficiently gather the metrics it needs with minimal complexity. This approach leverages existing AWS services in a simple and direct manner, aligning well with the company's requirements.

upvoted 3 times

  **DevopsCircus** 1 year, 4 months ago

Answer A looks pretty good except one thing

"Configure a CloudWatch Logs metric filter that increments a metric for each API operation name."

Can the metric filter "increment" the metric or just send metric value?

Answer B looks weird to me - I don't know the way to populate Logs Insights query result as the CW metric

upvoted 1 times

🗨️ 👤 **zain1258** 1 year, 4 months ago

Selected Answer: A

The right option is A. In the question, the requirement is to get a metric. In option B, we are not creating any metric.

upvoted 1 times

🗨️ 👤 **sivre** 1 year, 5 months ago

For me its B. From the question "The company needs to gather a metric for each API operation by response code for each version of the application that is in use". Only CloudWatch Logs Insight is querying for response code and application version. CloudWatch Logs metric filter is using API operation name to create an aggregate metric

upvoted 1 times

🗨️ 👤 **rbm2023** 1 year, 5 months ago

Selected Answer: A

It should be between A and B but the thing is the difference use cases between CW Logs Metric and CW Logs Insights. Metrics will provide Aggregated data while insights give you direct access to raw log event data. Hence I would go for A

upvoted 2 times

A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization, the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.

Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.

A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.

Which solution will meet these requirements?

- A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.
- B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.
- C. Configure provisioned concurrency on the Lambda function. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
- D. Configure reserved concurrency on the Lambda function. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Suggested Answer: C

Community vote distribution

C (100%)

 **Saga**  5 months, 2 weeks ago

Selected Answer: C

To reduce the latency of the Lambda function at all times of the day, the best solution is to configure provisioned concurrency on the Lambda function with a concurrency value of 1 and also configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100 (Option C).

Provisioned concurrency will ensure that the Lambda function has a set number of instances always available, which will reduce the cold start time. By setting the provisioned concurrency values to a minimum of 1 and a maximum of 100, the Lambda function can handle sudden spikes in traffic and can scale down during low-traffic periods, thus minimizing costs.

upvoted 9 times

 **ele**  5 months, 2 weeks ago


Selected Answer: C

Lambda integrates with Application Auto Scaling, allowing you to manage provisioned concurrency on a schedule or based on utilization.

<https://docs.aws.amazon.com/lambda/latest/dg/provisioned-concurrency.html>

<https://docs.aws.amazon.com/autoscaling/application/userguide/services-that-can-integrate-lambda.html>

upvoted 2 times

 **c3518fc** 5 months, 2 weeks ago

Selected Answer: C

By implementing provisioned concurrency with auto-scaling and retaining the DynamoDB DAX cluster, the DevOps engineer can effectively reduce the latency of the Lambda function at all times of the day while ensuring that the application can handle varying request volumes.

upvoted 2 times

 **Gomer** 5 months, 2 weeks ago

Selected Answer: C

Answer is to use application autoscaling to create a Lambda scaling policy for Provisioned Concurrency based on a re-occurring schedule. Here is reference that explains exactly how to do it CLI: <https://aws.amazon.com/blogs/compute/scheduling-aws-lambda-provisioned-concurrency-for-recurring-peak-usage/>

Here are truncated command examples from the reference:

```
aws application-autoscaling register-scalable-target --service-namespace lambda [...] --min-capacity 1 --max-capacity 100 --scalable-dimension lambda:function:ProvisionedConcurrency
```

```
aws application-autoscaling put-scheduled-action --service-namespace lambda --scalable-dimension lambda:function:ProvisionedConcurrency --scalable-target-action MinCapacity=100 [...]
```

upvoted 1 times

🗨️ **omankoman** 6 months, 3 weeks ago

Selected Answer: C

C is right answer.

upvoted 1 times

🗨️ **NagaoShingo** 6 months, 4 weeks ago

Selected Answer: C

C is right answer. Omamko.

upvoted 1 times

🗨️ **thanhnv142** 1 year, 1 month ago

C definitely

upvoted 1 times

🗨️ **yuliaqwerty** 1 year, 2 months ago

Agree answer C. Provisioned concurrency – This is the number of pre-initialized execution environments allocated to your function. These execution environments are ready to respond immediately to incoming function requests.

upvoted 1 times

🗨️ **zijo** 1 year, 3 months ago

Auto Scaling makes it easy to dynamically adjust the provisioned concurrency based on metrics and hence option C is a good choice to dynamically adjust based on the changing demand levels of the lambda function throughout the day.

upvoted 1 times

🗨️ **SanChan** 1 year, 9 months ago

C, This can help to optimize the number of instances available to serve requests, reducing the likelihood of cold starts and improving overall performance.

upvoted 2 times

🗨️ **alce2020** 1 year, 11 months ago

C it is

upvoted 1 times

🗨️ **jqso234** 1 year, 11 months ago

C is correct

upvoted 1 times

🗨️ **lqp0_0qpl** 1 year, 11 months ago

D / C - Provisioned concurrency is a manually set fixed value.

upvoted 1 times

🗨️ **Gomer** 9 months, 3 weeks ago

Somewhat true, except (in researching this), I discovered it the setting can also be dynamically adjusted up or down by application auto-scaling based on a schedule. Look at my other references, and you'll understand.

upvoted 1 times

A company is adopting AWS CodeDeploy to automate its application deployments for a Java-Apache Tomcat application with an Apache Webserver. The development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production. The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that they can set different log level configurations depending on the deployment group without having a different application revision for each group.

How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?


- A. Tag the Amazon EC2 instances depending on the deployment group. Then place a script into the application revision that calls the metadata service and the EC2 API to identify which deployment group the instance is part of. Use this information to configure the log level settings. Reference the script as part of the AfterInstall lifecycle hook in the appspec.yml file.
- B. Create a script that uses the CodeDeploy environment variable `DEPLOYMENT_GROUP_NAME` to identify which deployment group the instance is part of. Use this information to configure the log level settings. Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file.
- C. Create a CodeDeploy custom environment variable for each environment. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part of. Use this information to configure the log level settings. Reference this script as part of the ValidateService lifecycle hook in the appspec.yml file.
- D. Create a script that uses the CodeDeploy environment variable `DEPLOYMENT_GROUP_ID` to identify which deployment group the instance is part of to configure the log level settings. Reference this script as part of the Install lifecycle hook in the appspec.yml file.

Suggested Answer: B

Community vote distribution

B (88%)

12%

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

B. In the docs: <https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html> you'll find a Note: "The Start, Install, TestTraffic, AllowTraffic, and End events in the deployment cannot be scripted, which is why they appear in gray in this diagram." That's why it's not D.

upvoted 17 times

 **y0eri** Highly Voted 5 months, 2 weeks ago

Answer: B

Read this blog.

<https://aws.amazon.com/blogs/devops/using-codedeploy-environment-variables/>

```
if [ "$DEPLOYMENT_GROUP_NAME" == "Staging" ]
then
sed -i -e 's/LogLevel warn/LogLevel debug/g' /etc/httpd/conf/httpd.conf
fi
```

upvoted 11 times

 **steli0** Most Recent 3 months, 2 weeks ago

Selected Answer: B

On top of what others wrote the other difference between B and D is the BeforeInstall, since you need to configure log level before deploying the code/service.

upvoted 1 times

 **Sazeka** 5 months, 2 weeks ago

Selected Answer: B

B is correct.

`DEPLOYMENT_ID` : This variable contains the deployment ID of the current deployment.

`DEPLOYMENT_GROUP_NAME` : This variable contains the name of the deployment group. A deployment group is a set of instances associated with an application that you target for a deployment.

upvoted 2 times

🗨️ 👤 **zijo** 5 months, 2 weeks ago

Answer B. You only need to consider options B and D which are the least complex ones. The option B gives the CodeDeploy environment variable `DEPLOYMENT_GROUP_NAME` that points to different instances and gives the option to set different log-level configurations in the same script depending on the deployment group without having a different application revision for each group. Also, The `BeforeInstall` lifecycle hook in the `appspec.yml` file refers to a script that will run on the instance before the application revision files are installed.

upvoted 2 times

🗨️ 👤 **c3518fc** 5 months, 2 weeks ago

Selected Answer: B

version: 0.0

os: linux

files:

- source: /

destination: /var/www/html/

hooks:

BeforeInstall:

- location: scripts/update_log_level.sh

timeout: 300

runas: root

upvoted 1 times

🗨️ 👤 **Gomer** 5 months, 2 weeks ago

Selected Answer: B

This reference specifies the exact scenario described in "B", so I have to go with that "Set the log level according to the deployment group."

<https://aws.amazon.com/blogs/devops/using-codedeploy-environment-variables/>

```
cat install_dependencies.sh
```

```
[...]
```

```
if [ "$DEPLOYMENT_GROUP_NAME" == "Staging" ]
```

```
[...]
```

```
cat appspec.yml
```

```
hooks:
```

```
BeforeInstall:
```

```
- location: scripts/install_dependencies
```

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

B is correct

upvoted 2 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

B is correct: <without having a different application revision for each group> means A and C is incorrect.

A and C: <place a script into the application revision> both mention this, indicating a revision of the app, which is contradicted to the question

D: there is no Install lifecycle hook

upvoted 1 times

🗨️ 👤 **hoakhanh281** 1 year, 2 months ago

Selected Answer: D

Answer D.

You only need to consider options B and D which are the least complex ones. The option B gives the CodeDeploy environment variable `DEPLOYMENT_GROUP_NAME`, but `DEPLOYMENT_GROUP_ID` is recommended because it's more reliable and less prone to changes or inconsistencies. Answer D with settings `DEPLOYMENT_GROUP_ID` environment variable, which contains the unique identifier for the deployment group. This allows you to identify the deployment group without relying on custom scripts or metadata services.

upvoted 3 times

🗨️ 👤 **wem** 1 year, 3 months ago

Based on the analysis, Option B is the most efficient and straightforward approach. It uses the built-in `DEPLOYMENT_GROUP_NAME` environment variable provided by CodeDeploy and involves minimal management overhead. The script can easily read this variable to determine the

deployment group and set the log level accordingly, eliminating the need for different script versions for each group. This method aligns well with the requirement of least management overhead and simplicity.

upvoted 1 times

🗨️ 👤 **SanChan** 1 year, 9 months ago

B is the most straightforward and efficient solution to meet the requirements with the least management overhead and without requiring different script versions for each deployment group.

upvoted 2 times

🗨️ 👤 **Aja1** 1 year, 7 months ago

<https://aws.amazon.com/blogs/devops/using-codedeploy-environment-variables/>

upvoted 1 times

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

Answer is B practical use case

upvoted 1 times

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

Selected Answer: B

Running the hook during the Install or AfterInstall would make more sense but hooks for Install are not available (like in answer D) and AfterInstall is not included in answers so the best answer is B.

upvoted 2 times

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

Selected Answer: B

B is the only correct answer

upvoted 4 times

🗨️ 👤 **alce2020** 1 year, 11 months ago

B. Create a script that uses the CodeDeploy environment variable `DEPLOYMENT_GROUP_NAME` to identify which deployment group the instance is part of, and use this information to configure the log level settings. Reference this script as part of the `BeforeInstall` lifecycle hook in the `appspec.yml` file, would be the option with the least management overhead and without requiring different script versions for each deployment group

upvoted 3 times

🗨️ 👤 **jqso234** 1 year, 11 months ago

Option B is the best solution for this use case. By using the CodeDeploy environment variable `DEPLOYMENT_GROUP_NAME`, the script can identify the deployment group that the instance is part of, without requiring any additional configuration or management overhead. The script can then dynamically configure the log level settings based on the identified deployment group.

upvoted 1 times

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

B See: <https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html>

upvoted 1 times

A company requires its developers to tag all Amazon Elastic Block Store (Amazon EBS) volumes in an account to indicate a desired backup frequency. This requirement includes EBS volumes that do not require backups. The company uses custom tags named Backup_Frequency that have values of none, daily, or weekly that correspond to the desired backup frequency. An audit finds that developers are occasionally not tagging the EBS volumes.

A DevOps engineer needs to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified.



Which solution will meet these requirements?

- A. Set up AWS Config in the account. Create a custom rule that returns a compliance failure for all Amazon EC2 resources that do not have a Backup Frequency tag applied. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- B. Set up AWS Config in the account. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- C. Turn on AWS CloudTrail in the account. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly. Specify the runbook as the target of the rule.
- D. Turn on AWS CloudTrail in the account. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume events. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly. Specify the runbook as the target of the rule.

Suggested Answer: B

Community vote distribution

B (100%)

  **bb4f13b** 2 months, 1 week ago

Selected Answer: D

Option D:

Real-Time Enforcement: Ensures tagging compliance as soon as a volume is created or modified.

Comprehensive Coverage: Captures both CreateVolume and ModifyVolume events.

Minimal Overhead: Automates tagging without requiring manual audits or remediation actions

upvoted 1 times

  **dark4igi** 1 week, 5 days ago

what will happen if dev decide to change tag to daily?

AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly

upvoted 1 times

  **lunt** 5 months, 2 weeks ago

Selected Answer: B

Only takes few minutes to login > Config > Managed rulename = BACKUP_PLAN_MIN_FREQUENCY_AND_MIN_RETENTION_CHECK

A = tags everything in EC2, thats EC2:* which includes ELB/EIP/etc. Nope.

Option B you can specify the tags to match & expected values = answer.

upvoted 3 times

  **CristianoRosa** 5 months, 2 weeks ago

Selected Answer: B

A: It works, but it uses a custom rule.

B: It is simpler than option A as it uses a managed rule which already exists.

C: It only applies to new volumes and does not address existing resources.

D: It is better than C but still does not fully meet the requirement to check all EBS volumes and enforce compliance.

Best Answer is B.



upvoted 3 times

  **c3518fc** 5 months, 2 weeks ago

Selected Answer: B

By leveraging the AWS Config managed rule and automated remediation action, the DevOps engineer can ensure that all EBS volumes in the account always have the required Backup_Frequency tag, enabling the company to perform backups at least weekly unless a different value is explicitly specified. This solution provides continuous monitoring and automated remediation, reducing the risk of human error and ensuring compliance with the company's backup policy.

upvoted 1 times

  **ajeeshb** 5 months, 2 weeks ago



Selected Answer: B

Option B --> AWS config managed rule on EC2::Volume resource + custom SSM automation document

Not Option A --> because it says custom config rule on all EC2::Instance + Managed SSM automation document

Not options C & D --> As it says cloudtrail which is for logging API actions

upvoted 2 times

  **ajeeshb** 8 months, 1 week ago

sorry, a typo.. Option A also says custom SSM automation document, but it is wrong where it says custom config rule on all Ec2::Instance

upvoted 1 times

  **Diego1414** 1 year ago

Answer is A.

Checks if your resources have the tags that you specify. For example, you can check whether your Amazon EC2 instances have the CostCenter tag, while also checking if all your RDS instance have one set of Keys tag. Separate multiple values with commas. You can check up to 6 tags at a time.

The AWS-managed AWS Systems Manager automation document AWS-SetRequiredTags does not work as a remediation with this rule. You will need to create your own custom Systems Manager automation documentation for remediation

<https://docs.aws.amazon.com/config/latest/developerguide/required-tags.html>

upvoted 1 times

  **Hizumi** 1 year ago

We don't need to create a custom AWS Config rule, we can utilize the managed rule to detect for non-compliance on the EBS volumes.

Otherwise the options indicate to use a custom runbook for AWS Systems Manager to remediate the missing tags.

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

B is correct: We should use AWS config for this task

C and D: cloud trail is for auditing account activities, which is irrelevant

A: <returns a compliance failure for all Amazon EC2 resources> : we need to remediate EC2 volumes only, not all EC2 resources

upvoted 4 times

  **Sisanda_giiven** 1 year, 1 month ago

A is the correct answer "The AWS-managed AWS Systems Manager automation document AWS-SetRequiredTags does not work as a remediation with this rule. You will need to create your own custom Systems Manager automation documentation for remediation." from this link :

<https://docs.aws.amazon.com/config/latest/developerguide/required-tags.html>

upvoted 1 times

  **zijo** 1 year, 3 months ago

B is the best choice. If you look at Config Managed Rules you can find - ebs-in-backup-plan - Check if Amazon Elastic Block Store (Amazon EBS) volumes are added in backup plans of AWS Backup. The rule is NON_COMPLIANT if Amazon EBS volumes are not included in backup plans.

upvoted 3 times

  **SanChan** 1 year, 9 months ago

Selected Answer: B

B is the most straightforward and efficient solution to ensure that all EBS volumes always have the Backup_Frequency tag applied with the least amount of effort.

A This approach requires more effort than using a managed rule provided by AWS.

upvoted 3 times

  **Aja1** 1 year, 7 months ago

<https://docs.aws.amazon.com/config/latest/developerguide/required-tags.html>

upvoted 4 times

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

Selected Answer: B

B makes sense, you can use managed rule "required-tags" to identify non-compliant volumes and custom SSM document to fix it.
upvoted 1 times

🗳️ 👤 **vhernan** 1 year, 10 months ago

Selected Answer: B

B makes sense
upvoted 1 times

🗳️ 👤 **alice2020** 1 year, 11 months ago

B. Set up AWS Config in the account. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
upvoted 1 times

🗳️ 👤 **ele** 1 year, 11 months ago

Selected Answer: B

Answer B: Config has a managed rule for type AWS EC2 Volume for tag compliance check.
upvoted 1 times

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

Selected Answer: B

B for me. <https://aws.amazon.com/ru/blogs/mt/build-an-aws-config-custom-rule-to-optimize-amazon-ebs-volume-types/>
upvoted 2 times

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

Sorry A is the answer. This is custom rule
upvoted 2 times

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

But very strange that custom rule for all ec2 instances , it should be only ec2 volumes
upvoted 1 times

🗳️ 👤 **jqso234** 1 year, 11 months ago

Option A creates a custom rule that applies to all EC2 resources, not just volumes, which may create additional overhead. The custom AWS Systems Manager Automation runbook is used to apply the Backup_Frequency tag with a value of weekly, but this approach can result in inconsistent tagging if the developers specify a different desired backup frequency. Therefore, Option A is not the correct answer.

Option B is the correct answer because it uses a managed rule specifically for EC2 volumes, which simplifies the configuration effort and ensures that all volumes have the Backup_Frequency tag applied consistently. The custom AWS Systems Manager Automation runbook is used to automatically apply the Backup_Frequency tag with a value of weekly, which reduces the risk of data loss due to missing backups. Your comment that the managed rule should only apply to volumes is correct, and Option B addresses that requirement.


upvoted 4 times

A company is using an Amazon Aurora cluster as the data store for its application. The Aurora cluster is configured with a single DB instance. The application performs read and write operations on the database by using the cluster's instance endpoint. The company has scheduled an update to be applied to the cluster during an upcoming maintenance window. The cluster must remain available with the least possible interruption during the maintenance window. What should a DevOps engineer do to meet these requirements?

- A. Add a reader instance to the Aurora cluster. Update the application to use the Aurora cluster endpoint for write operations. Update the Aurora cluster's reader endpoint for reads.
- B. Add a reader instance to the Aurora cluster. Create a custom ANY endpoint for the cluster. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.
- C. Turn on the Multi-AZ option on the Aurora cluster. Update the application to use the Aurora cluster endpoint for write operations. Update the Aurora cluster's reader endpoint for reads.
- D. Turn on the Multi-AZ option on the Aurora cluster. Create a custom ANY endpoint for the cluster. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.

Suggested Answer: C

Community vote distribution



junrun3 Highly Voted 1 year, 8 months ago

Selected Answer: A

B and D are incorrect because Aurora cluster provides cluster and read endpoints, but does not support creating custom ANY endpoints.

C and D are incorrect because Amazon Aurora's multi-AZ option must be set when the DB instance is created.

Therefore, A is correct.

upvoted 15 times

Just_Ninja Highly Voted 1 year, 8 months ago

Option A is the right choice for an existing Cluster without Multi-AZ!

Refer to: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Concepts.AuroraHighAvailability.html>

#Read the Tip Box#

"You can set up a Multi-AZ cluster by making a simple choice when you create the cluster. The choice is simple whether you use the AWS Management Console, the AWS CLI, or the Amazon RDS API. You can also make an existing Aurora cluster into a Multi-AZ cluster by adding a new reader instance and specifying a different Availability Zone."

upvoted 12 times

haazybanj Most Recent 5 months, 2 weeks ago

Selected Answer: C

C. Turn on the Multi-AZ option on the Aurora cluster. Update the application to use the Aurora cluster endpoint for write operations. Update the Aurora cluster's reader endpoint for reads.

Enabling Multi-AZ ensures that the data in the Aurora cluster is replicated across multiple Availability Zones (AZs), providing high availability and durability. During maintenance, the update will be applied to one AZ at a time, allowing the cluster to remain available. Updating the application to use the cluster endpoint for write operations ensures that writes will continue to be directed to the primary instance in the cluster, while updating the reader endpoint for reads allows read traffic to be routed to the appropriate instance. Adding a reader instance or creating a custom ANY endpoint are not necessary for meeting the requirement of minimizing interruption during maintenance.

upvoted 2 times



koenigParas2324 5 months, 2 weeks ago

Selected Answer: C

To meet the requirement of keeping the Aurora cluster available with the least possible interruption during the maintenance window, the DevOps engineer should choose option C: Turn on the Multi-AZ option on the Aurora cluster. Update the application to use the Aurora cluster endpoint for write operations. Update the Aurora cluster's reader endpoint for reads.

Option A is incorrect because adding a reader instance alone does not provide high availability during maintenance, and updating the application to use the reader endpoint for reads is unnecessary when the Multi-AZ option is enabled

upvoted 1 times



  **alexleely** 5 months, 2 weeks ago

Selected Answer: A

Option A is the correct choice, adding a reader instance after provisioning is the same as setting a Multi-AZ during creation. In the event that the primary fails, the reader instance will be promoted to do both reading and writing automatically.

Additionally, reader instance can also be used for read activity if you use the cluster reader endpoint which you can serve to user/application closer to the region for better performance.

upvoted 2 times

  **vietnguyen2** 5 months, 2 weeks ago

Selected Answer: A

"You can set up a Multi-AZ DB cluster by making a simple choice when you create the cluster. You can use the AWS Management Console, the AWS CLI, or the Amazon RDS API. You can also convert an existing Aurora DB cluster into a Multi-AZ DB cluster by adding a new reader DB instance and specifying a different Availability Zone."

upvoted 1 times

  **Rizwan_Shaukat** 5 months, 2 weeks ago

Selected Answer: A

To meet the requirements of the given scenario, the DevOps engineer should do the following:

Add a reader instance to the Aurora cluster:

This will allow the application to offload read operations to the reader instance, reducing the load on the primary instance.

The application should be updated to use the Aurora cluster endpoint for write operations and the reader endpoint for read operations.



The engineer should not turn on the Multi-AZ option on the Aurora cluster.

Multi-AZ is used to provide high availability and failover capabilities, but it does not necessarily minimize interruption during a maintenance window.

Adding a reader instance is a more appropriate solution to maintain availability and distribute the read workload.

Therefore, the correct option is A. Add a reader instance to the Aurora cluster and update the application to use the appropriate endpoints for read and write operations.


upvoted 2 times

  **Rahul369** 8 months, 3 weeks ago

Selected Answer: A

You cannot change the az option after creation but can deploy a reader instance in another az and use it for reading and writing in your main instance.

upvoted 1 times

  **Malcnorth59** 8 months, 3 weeks ago

There isn't a specific Multi-AZ mode for Aurora it's multi-AZ by default (it uses all three AZs). So I think A or B. A is the most commonly used method but I think B offers less disruption because requests are always routed to the instance with greater availability.



upvoted 1 times

  **c3518fc** 10 months ago

Selected Answer: D

By enabling Multi-AZ deployment, creating a custom ANY endpoint, and updating the application to use this endpoint for all read and write operations, the DevOps engineer can ensure that the Aurora cluster remains available during the maintenance window with minimal interruption. The application will be able to transparently connect to the available instances (primary or read-only replica), and Aurora will automatically fail over to the read-only replica if the primary instance becomes unavailable during the maintenance process.

upvoted 1 times

  **01037** 10 months, 1 week ago

Selected Answer: B

Can't tell the difference between A and C except Multi-AZ, both should be working if only read is needed during maintenance window.

And also not understand why only read is needed when people choose them.

Some say there is no custom ANY endpoint, I think it only means you can choose any instance or instances to that endpoint.

So I go with B

upvoted 1 times

🗨️ 👤 **Mackn** 11 months, 3 weeks ago

answer A).

the question doesn't mention what DB is behind the Aurora.

Multi-AZ config avoids downtime EXCEPT MySQL/MariaDB.

So the question mentions "the least possible interruption", then A) is the appropriate one

upvoted 1 times

🗨️ 👤 **kyuhuck** 1 year ago

is corret is 'c' - > 'a' is not

This option leverages Aurora's built-in high availability and failover mechanisms to ensure minimal interruption. By using the cluster endpoint for writes, the application automatically writes to the primary instance. In case of maintenance or failure, Aurora handles failover to another instance with minimal downtime. The reader endpoint distributes read traffic across available replicas, enhancing read scalability and availability without affecting write operations. This setup ensures that the application remains as available as possible during maintenance

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

A is good

upvoted 1 times

🗨️ 👤 **Jonalb** 1 year, 1 month ago

A is correct.

upvoted 1 times

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

Agree answer A. There is no ANY custom endpoint and multy-AZ can be set up during cluster creation

upvoted 1 times

🗨️ 👤 **n00b2023** 1 year, 2 months ago

'A' - since Multi AZ has to be setup when the cluster is created. It cannot be updated later.

upvoted 2 times

🗨️ 👤 **phu0298** 11 months, 3 weeks ago

I agree with you. C D is the wrong selection. because we can't enable Multi-AZ after cluster is created

upvoted 1 times

A company must encrypt all AMIs that the company shares across accounts. A DevOps engineer has access to a source account where an unencrypted custom AMI has been built. The DevOps engineer also has access to a target account where an Amazon EC2 Auto Scaling group will launch EC2 instances from the AMI. The DevOps engineer must share the AMI with the target account. The company has created an AWS Key Management Service (AWS KMS) key in the source account. Which additional steps should the DevOps engineer perform to meet the requirements? (Choose three.)


- A. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.
- B. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.
- C. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
- D. In the source account, modify the key policy to give the target account permissions to create a grant. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.
- E. In the source account, share the unencrypted AMI with the target account.
- F. In the source account, share the encrypted AMI with the target account.

Suggested Answer: ACD

Community vote distribution

ADF (96%)

2%

 **kacsabacs78** Highly Voted 1 year, 9 months ago

Selected Answer: ADF

ADF seems to be the correct answer
upvoted 11 times

 **Dimidrol** Highly Voted 1 year, 11 months ago

Selected Answer: ADF

A D F for me. <https://jackiechen.blog/2020/01/29/share-encrypted-ami-across-aws-accounts/>
upvoted 9 times

 **cmzk** Most Recent 2 weeks, 5 days ago

Selected Answer: ADF

A. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.

The DevOps engineer must first encrypt the unencrypted AMI using the KMS key created in the source account. This creates a new encrypted AMI.

D. In the source account, modify the key policy to give the target account permissions to create a grant. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.

The KMS key policy in the source account must be updated to allow the target account to use the key for decryption. Additionally, the target account needs to create a KMS grant to allow the Auto Scaling group service-linked role to use the KMS key.

F. In the source account, share the encrypted AMI with the target account.

After encrypting the AMI, the DevOps engineer must share the encrypted AMI with the target account by modifying the AMI's permissions.
upvoted 1 times

 **cmzk** 2 weeks, 5 days ago

Incorrect Choices:

B. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.

This is incorrect because the company requires the use of the KMS key created in the source account, not the default EBS encryption key.

C. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.

This is incorrect because KMS grants cannot be created for resources in another account. The target account must create the grant for its own service-linked role.

E. In the source account, share the unencrypted AMI with the target account.

This is incorrect because the company requires all AMIs to be encrypted before sharing.

upvoted 1 times

🗲️ 👤 **cmzk** 2 weeks, 5 days ago

Selected Answer: ADF

A, D, F are the correct

upvoted 1 times

🗲️ 👤 **namtp** 7 months, 1 week ago

Selected Answer: ADF

ADF for me,

upvoted 1 times

🗲️ 👤 **martinarg2024** 1 year ago

Selected Answer: ADF

ADF is correct

upvoted 1 times

🗲️ 👤 **Vitalydt** 1 year ago

Selected Answer: ADF

A D F for me

upvoted 1 times

🗲️ 👤 **thanhnv142** 1 year, 1 month ago

ADF:

A: cannot be B because using KMS

D: Must share with the account because grant is only temp

F: share the AMI with the target

upvoted 2 times

🗲️ 👤 **thanhnv142** 1 year, 1 month ago

AFD seem about right

upvoted 1 times

🗲️ 👤 **Jonalb** 1 year, 1 month ago

ADF the correct answer

upvoted 1 times

🗲️ 👤 **khchan123** 1 year, 1 month ago

Selected Answer: ACF

ACF. For autoscaling to work a KMS grant is needed

upvoted 1 times

🗲️ 👤 **khchan123** 1 year, 1 month ago

Should be ADF

upvoted 1 times

🗲️ 👤 **harithzainudin** 1 year, 3 months ago

Selected Answer: ADF

ADF is the right answer

upvoted 1 times

🗲️ 👤 **VrilianVirgil** 1 year, 4 months ago

Selected Answer: ADF

C is incorrect as the AMI ****MUST**** be shared with the account.

not just the scaling group. So it would make sense for the target account to create the grant.

upvoted 2 times

🗲️ 👤 **ataince** 1 year, 5 months ago

Selected Answer: ADF

ADF is the right answer.

upvoted 1 times

🗲️ 👤 **BaburTurk** 1 year, 6 months ago

Selected Answer: ADF

<https://aws.amazon.com/blogs/security/how-to-create-a-custom-ami-with-encrypted-amazon-ebs-snapshots-and-share-it-with-other-accounts-and-regions/>

upvoted 3 times

🗲️ 👤 **Skshitiz** 1 year, 6 months ago

Selected Answer: ADF

ADF is right

upvoted 1 times

🗲️ 👤 **DavidPham** 1 year, 7 months ago

Selected Answer: ADF

ADF is correct

upvoted 2 times

A company uses AWS CodePipeline pipelines to automate releases of its application. A typical pipeline consists of three stages: build, test, and deployment. The company has been using a separate AWS CodeBuild project to run scripts for each stage. However, the company now wants to use AWS CodeDeploy to handle the deployment stage of the pipelines.

The company has packaged the application as an RPM package and must deploy the application to a fleet of Amazon EC2 instances. The EC2 instances are in an EC2 Auto Scaling group and are launched from a common AMI.

Which combination of steps should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Create a new version of the common AMI with the CodeDeploy agent installed. Update the IAM role of the EC2 instances to allow access to CodeDeploy.
- B. Create a new version of the common AMI with the CodeDeploy agent installed. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.
- C. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the Auto Scaling group as the deployment target. Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AMI. Configure CodeDeploy to deploy the newly created AMI.
- D. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the Auto Scaling group as the deployment target. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.
- E. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the EC2 instances that are launched from the common AMI as the deployment target. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.

Suggested Answer: AD

Community vote distribution

AD (100%)

  **bcx** Highly Voted 1 year, 9 months ago



A and D are the correct ones.

E is wrong because it says that the instances are on an ASG.

C is wrong. You deploy the new RPM on the AMI, you do not create a new AMI every time to install the RPM.

B is wrong, the appspec has nothing to do with permissions

upvoted 10 times

  **thanhnv142** Highly Voted 5 months, 2 weeks ago

Selected Answer: AD

A and D are correct:

A: rebuild the AMI and Update the IAM role of the EC2 instances to allow access to CodeDeploy is necessary

B: no need to grant AppSpec file access to code CodeDeploy

C: <Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AMI> : this is unnecessary, we have already done it in option A. Additionally, recreating AMI each time running the CI/CD pipeline is unnecessary

D: ok

E: <Specify the EC2 instances that are launched from the common AMI as the deployment target>: this is time-consuming. There might be hundreds of EC2 instances and targeting them individually is time-consuming and not effective.

upvoted 6 times

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

the ans is A and D

D not E because in deployment the best practise deploy in group not instance

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

A and D:

B is incorrect: AppSpec file does not need to be granted access to code deploy. It is code deploy that needs the permission to get access to Appspec file

upvoted 1 times

🗳️ 👤 **z_inderjot** 1 year, 2 months ago

Selected Answer: AD

A - instances need code deploy agent and role

D - target as ASG

upvoted 3 times

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

A D.

IAM role/instance profile requirement for EC2 is to allow EC2 access to S3 buckets used by CodeDeploy.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/getting-started-create-iam-instance-profile.html>

upvoted 2 times

🗳️ 👤 **DZ_Ben** 1 year, 4 months ago

Should be BD. EC2 doesn't need a permission to access CodeDeploy. Instead an IAM role associated with Code Deployment Group should have an permission to launch instances in that autoscaling group.

upvoted 2 times

🗳️ 👤 **z_inderjot** 1 year, 2 months ago

A is right

you have to attach IAM role to EC2 instance , for them to be controlled by Code deploy . The agent running in the ec2 instance needs to talk with code deploy .

upvoted 3 times

🗳️ 👤 **sivre** 1 year, 5 months ago

Why EC2 instance need access to CodeDeploy??, in the doc is mentioned only S3: "Create or locate an IAM instance profile that allows the Amazon EC2 Auto Scaling group to work with Amazon S3" <https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

upvoted 3 times

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

As the codedeploy agent is being installed inside the EC2, This agent facilitates the deployment process by coordinating with the CodeDeploy service. Hence,

The EC2 instances must have an IAM role that grants them the necessary permissions to interact with CodeDeploy. This step is critical to ensure that the deployment process can be executed securely and successfully

upvoted 3 times

🗳️ 👤 **Ja13** 1 year, 7 months ago

Selected Answer: AD

AD as explained en the comments

upvoted 1 times

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

Selected Answer: AD

AD are correct.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

upvoted 2 times

🗳️ 👤 **Aja1** 1 year, 7 months ago

An in-place deployment allows you to deploy your application without creating new infrastructure.

The deployment type that is specific to the deployment's compute platform or deployments initiated by a CloudFormation stack

update.<https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

upvoted 2 times

🗳️ 👤 **tycho** 1 year, 10 months ago

A, D; B is wrong, because the AppSpec file cannot grant permissions for CodeDeploy

upvoted 1 times

🗳️ 👤 **alce2020** 1 year, 11 months ago

A&D it is

upvoted 1 times

🗳️ 👤 **ele** 1 year, 11 months ago

Selected Answer: AD

A,D.

B - wrong, the 'permissions' section specifies how special permissions, should be applied to the files and directories/folders in the 'files' section

C wrong, no need as AMI was already built.

E wrong, as ASG is the target.

upvoted 2 times


A company's security team requires that all external Application Load Balancers (ALBs) and Amazon API Gateway APIs are associated with AWS WAF web ACLs. The company has hundreds of AWS accounts, all of which are included in a single organization in AWS Organizations. The company has configured AWS Config for the organization. During an audit, the company finds some externally facing ALBs that are not associated with AWS WAF web ACLs.

Which combination of steps should a DevOps engineer take to prevent future violations? (Choose two.)

- A. Delegate AWS Firewall Manager to a security account.
- B. Delegate Amazon GuardDuty to a security account.
- C. Create an AWS Firewall Manager policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
- D. Create an Amazon GuardDuty policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
- E. Configure an AWS Config managed rule to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

Suggested Answer: AC

Community vote distribution



ataince Highly Voted 1 year, 5 months ago

Selected Answer: AC

If you see WAF you have to think AWS Firewall Manager.
upvoted 11 times

alce2020 Highly Voted 1 year, 11 months ago

A and C
upvoted 9 times

ele Most Recent 5 months, 2 weeks ago

Selected Answer: AC

If instead you want to automatically apply the policy to existing in-scope resources, choose Auto remediate any noncompliant resources. This option creates a web ACL in each applicable account within the AWS organization and associates the web ACL with the resources in the accounts.

When you choose Auto remediate any noncompliant resources, you can also choose to remove existing web ACL associations from in-scope resources, for the web ACLs that aren't managed by another active Firewall Manager policy. If you choose this option, Firewall Manager first associates the policy's web ACL with the resources, and then removes the prior associations. If a resource has an association with another web ACL that's managed by a different active Firewall Manager policy, this choice doesn't affect that association.

upvoted 1 times

namtp 7 months, 1 week ago

Selected Answer: AC

I think that is best way to centralize manage firewall config
upvoted 1 times

jamesf 7 months, 2 weeks ago

Selected Answer: AC

As my understanding, WAF related with AWS Firewall Manager.
upvoted 1 times

Gomer 9 months, 3 weeks ago

Selected Answer: AC

These references indicate this can all be handled within Firewall manager (w/no references to Config or GuardDuty)

<https://aws.amazon.com/blogs/security/how-to-enforce-a-security-baseline-for-an-aws-waf-acl-across-your-organization-using-aws-firewall-manager/>

<https://aws.amazon.com/solutions/implementations/automations-for-aws-firewall-manager/>

upvoted 2 times

Gomer 9 months, 3 weeks ago

In reading a little further, I suspect that Config may be being used in the background (since Config must be enabled to use WAF. However, I believe that is totally transparent to the Organization WAF Administrator. The administration of WAF and enforcement of WAF policies is ALL handled with the Web Application Firewall service.

upvoted 1 times

🗲️ 👤 **01037** 10 months, 1 week ago

Selected Answer: AC

I think E works, but Firewall manager is designed for the purpose.

upvoted 1 times

🗲️ 👤 **Cervus18** 12 months ago

Selected Answer: AC

A and C: AWS Config rules are primarily used for monitoring and evaluating the configurations of your AWS resources for compliance with desired configurations. However, AWS Config also supports remediation actions through AWS Systems Manager Automation documents.

upvoted 1 times

🗲️ 👤 **Vitalydt** 1 year ago

Selected Answer: CE

Why not E?

upvoted 1 times

🗲️ 👤 **01037** 10 months, 1 week ago

I think E works, but Firewall manager is designed for the purpose.

upvoted 1 times

🗲️ 👤 **Cervus18** 12 months ago

AWS Config rules are primarily used for monitoring and evaluating the configurations of AWS resources for compliance with desired configurations. However, AWS Config also supports remediation actions through AWS Systems Manager Automation documents or lambda. Firewall manager is used to apply and enforce WebACLs to all ALBs at an organizational level to all your AWS Organization's accounts, and you can configure auto remediation for any non-compliant resource in any account.

upvoted 1 times

🗲️ 👤 **thanhnv142** 1 year, 1 month ago

A and C: Config does not have any action, only notifications

upvoted 2 times

🗲️ 👤 **Fco_Javier** 1 year, 7 months ago

A) is a prerequisites: AWS Firewall Manager prerequisites

https://docs.aws.amazon.com/es_es/waf/latest/developerguide/join-aws-orgs.html

upvoted 2 times

🗲️ 👤 **habros** 1 year, 8 months ago

Selected Answer: AC

GuardDuty only posts findings, hence they can be eliminated.

From my knowledge, Config only notifies.

Hence, A and C.

upvoted 2 times

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

Selected Answer: AC

A C for me

upvoted 1 times

A company uses AWS Key Management Service (AWS KMS) keys and manual key rotation to meet regulatory compliance requirements. The security team wants to be notified when any keys have not been rotated after 90 days. Which solution will accomplish this?

- A. Configure AWS KMS to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.
- B. Configure an Amazon EventBridge event to launch an AWS Lambda function to call the AWS Trusted Advisor API and publish to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. Develop an AWS Config custom rule that publishes to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.
- D. Configure AWS Security Hub to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **thanhnv142** Highly Voted 7 months, 2 weeks ago

C is correct

A is not because KMS does not provide this function

upvoted 5 times

🗳️ 👤 **yuliaqwerty** Most Recent 8 months, 1 week ago

Answer C. AWS Config

upvoted 3 times

🗳️ 👤 **habros** 1 year, 2 months ago

Selected Answer: C

C. Config rules notifies.

upvoted 3 times

🗳️ 👤 **Toptip** 1 year, 2 months ago

Selected Answer: C

Are these questions really came from DOP-C02?

upvoted 3 times

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

Selected Answer: C

C makes sense. it should be a custom rule. Rule "access-keys-rotated" checks for access keys, not KMS keys.

upvoted 2 times

🗳️ 👤 **alce2020** 1 year, 4 months ago

C it is

upvoted 1 times

🗳️ 👤 **ele** 1 year, 5 months ago

Selected Answer: C

custom config: C

upvoted 1 times

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

Selected Answer: C

Looks like C, actually there is a managed rule for this:

<https://docs.aws.amazon.com/config/latest/developerguide/access-keys-rotated.html>

anyway trusted advisor cannot be used as there is no such check, also KMS does not have this action, security hub is not conducting any active checks just react to events

upvoted 4 times

🗳️ 👤 **zijo** 6 months, 3 weeks ago

IAM Access Key & KMS key are different. The managed rule is for IAM Access key
upvoted 1 times

🗨️ 👤 **s50600822** 1 year, 1 month ago
access key?
upvoted 3 times

🗨️ 👤 **Dimidrol** 1 year, 5 months ago

Selected Answer: C

C for me. A there no such functionality, B i checked trusted advisor there is no such kms days, d is aggregator for config, guardduty. So you need config for D
upvoted 2 times

🗨️ 👤 **lqp0_0qpl** 1 year, 5 months ago
Tell me Why not D..
upvoted 1 times

🗨️ 👤 **beanxyz** 1 year ago

When you enable a control in Security hub it will automatically create a Config. There are 4 KMS related controls in security hub but none of them is about the rotation age. In this case you need to create a custom Config.
upvoted 1 times

🗨️ 👤 **Manny20** 1 year, 3 months ago

- Option D is not the correct answer because AWS Security Hub is primarily focused on aggregating and managing security findings, and it does not have a specific feature to monitor the age of AWS KMS keys.

upvoted 2 times

A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project. How can this issue be corrected in the MOST secure manner?

- A. Add the bucket name to the AllowedBuckets section of the CodeBuild project settings. Update the build spec to use the AWS CLI to download the database population script.
- B. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a token. Update the build spec to use cURL to pass the token and download the database population script.
- C. Remove unauthenticated access from the S3 bucket with a bucket policy. Modify the service role for the CodeBuild project to include Amazon S3 access. Use the AWS CLI to download the database population script.
- D. Remove unauthenticated access from the S3 bucket with a bucket policy. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **thanhnv142** Highly Voted 1 year, 1 month ago

C is correct:

- + Remove unauthenticated access from the S3 bucket with a bucket policy
 - + Modify the service role for the CodeBuild project to include Amazon S3 access.
- upvoted 5 times

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

Selected Answer: C

C is a correct answer.

Inside AWS, using of service roles is the best option.

upvoted 1 times

🗳️ 👤 **z_inderjot** 1 year, 2 months ago

Selected Answer: C

all these questions seem fairly to be part of aws devops exam

upvoted 3 times

🗳️ 👤 **zain1258** 1 year, 4 months ago

Selected Answer: C

C is correct

upvoted 1 times

🗳️ 👤 **Cervus18** 1 year, 4 months ago

Selected Answer: C

Involves using a service role also, which make it the most secure manner

upvoted 2 times

🗳️ 👤 **SanChan** 1 year, 9 months ago

Selected Answer: C

C is the correct answer because it involves removing unauthenticated access from the S3 bucket with a bucket policy, which ensures that only authorized users or services can access the bucket.

upvoted 4 times

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

Selected Answer: C



C is the best answer.

upvoted 1 times

🗳️ 👤 **alce2020** 1 year, 11 months ago

c is the answer

upvoted 2 times

  **ataince** 1 year, 11 months ago

c is the answer.

upvoted 1 times

  **ele** 1 year, 11 months ago

Selected Answer: C

C most secure

upvoted 1 times

An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0.

The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.

Which combination of steps will meet these requirements? (Choose three.)

- A. Create IAM policies that include the required permissions. Include the aws:PrincipalTag condition key.
- B. Create permission sets. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.
- C. Create a group in the IdP. Place users in the group. Assign the group to accounts and the permission sets in IAM Identity Center.
- D. Create a group in the IdP. Place users in the group. Assign the group to OUs and IAM policies.
- E. Enable attributes for access control in IAM Identity Center. Apply tags to users. Map the tags as key-value pairs.
- F. Enable attributes for access control in IAM Identity Center. Map attributes from the IdP as key-value pairs.

Suggested Answer: ABC

Community vote distribution



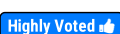
BCF (100%)

  **bcx**  1 year, 9 months ago

Selected Answer: BCF

I would go with BCF. I cannot make a large comment on why but manage an identity center setup at work and find that these are the correct ones IMHO. Your IdP has attributes, not tags, you have to rely on the IdP's attributes for instance. And you work with permission sets almost always, so the three answers about the permission sets make the full answer. You do not use IAM directly or tags for this.

upvoted 12 times

  **asfsdfsdf**  1 year, 11 months ago

Selected Answer: BCF

This is clearly stated here:

<https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/>

Answers are: BCF - permissions sets + IDP attributes mapping + groups

For example a user with IDP attribute of Dep/hr will be able to delete instances with this specific tag

upvoted 6 times

  **namtp**  7 months, 1 week ago

Selected Answer: BCF

BCF is correct answers.

Permission set + group created in the IdP, and map attributes is key

upvoted 1 times



  **Gomer** 9 months, 3 weeks ago

Selected Answer: BCF

While I have no great insights or expertise in this area, I do know how to read (RTFM) and quasi-solve the puzzle in my head. This reference URL (pdf) seems to touch all the steps listed in "B", "C", "F" and showed some extra steps not listed. Search and see for yourself.

https://d1.awsstatic.com/events/aws-reinforce-2022/IAM309_Designing-a-well-architected-identity-and-access-management-solution.pdf

upvoted 1 times

  **Gomer** 9 months, 3 weeks ago

Also, I might add, rather than just memorize the most votes answer to the question, I'd suggest actually going out to do some research and taking some long term notes you can reference later. That may take more time, but you also be more competent at work, and maybe keep your job longer. I love the fact that exam topics gives a forum to discuss and research complex questions and share findings. It's pretty lame If you come here to just memorize answers long enough to pass an exam.

upvoted 2 times

  **zijo** 1 year ago

Permission sets are stored in IAM Identity Center. So you know all answers that mention about permission sets and IAM Identity Center are likely correct

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

B, C, E seem more accurate:

B- need to attach the policy so that it can be usable. A is not true because IAM policies is not the same as in IAM Identity Center

C- not D because cannot assign group to IAM policies. IAM policies is attached to groups. also, need permission sets in Identity Center

E- attributes is basically tagging.

upvoted 1 times

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

correct answer seen as A-B-C. but 11 people sure the correct answer is B-C-F in discussion.

What is the answer?

Can the system show the correct answer as wrong or are people mistaken?

upvoted 1 times

🗨️ 👤 **davdan99** 1 year, 2 months ago

The examTopics answers in most cases are wrong, please read discussions, and references that users provide

upvoted 3 times

🗨️ 👤 **ajeeshb** 8 months, 1 week ago

Then why do people pay the fee for access, I dont understand. If it is from a discussion the people have to understand the answer (that too not very sure), why do they charge so much for the contributor access?!

upvoted 2 times

🗨️ 👤 **habros** 1 year, 8 months ago

<https://docs.aws.amazon.com/singlesignon/latest/userguide/provision-automatically.html>

upvoted 2 times

🗨️ 👤 **habros** 1 year, 8 months ago

Selected Answer: BCF

Example if I use IdP as my group, and I add users to the group, then my users will be onboarded via the SCIM method.

IAM roles does not apply to Control Tower landing zone. Hence B and C is secured (only permission sets for AWS SSO)

Does not make sense granting RBAC via tags...

upvoted 5 times

🗨️ 👤 **Aja1** 1 year, 7 months ago

An inline policy is a policy created for a single IAM identity (a user, group, or role). Inline policies maintain a strict one-to-one relationship between a policy and an identity

A permission set is a template that you create and maintain that defines a collection of one or more IAM policies.

upvoted 1 times

🗨️ 👤 **Aja1** 1 year, 7 months ago

IAM Identity Center helps you securely create, or connect, your workforce identities and manage their access centrally across AWS accounts and applications

Attribute mappings are used to map attribute types that exist in IAM Identity Center with like attributes in an AWS Managed Microsoft AD directory. IAM Identity Center retrieves user attributes from your Microsoft AD directory and maps them to IAM Identity Center user attributes. These IAM Identity Center user attribute mappings are also used for generating SAML assertions for your cloud applications.

upvoted 1 times

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

Selected Answer: BCF

BCF

<https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html>

upvoted 4 times

🗨️ 👤 **Rick365** 1 year, 9 months ago

Selected Answer: BCF

I beleive BCF

upvoted 1 times

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

Selected Answer: BCF

BCF makes more sense here.

upvoted 2 times

🗨️ 👤 **alice2020** 1 year, 11 months ago

ill go with B,C,F

upvoted 2 times

🗨️ 👤 **ele** 1 year, 11 months ago

Selected Answer: BCF

agree, BCF - permissions sets + IDP attributes mapping + groups

upvoted 2 times

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

A, C, E

upvoted 1 times

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (Amazon SQS) queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has auto scaling enabled for read and write capacity.


Which actions should a DevOps engineer take to resolve this delay? (Choose two.)

- A. Check the `ApproximateAgeOfOldestMessage` metric for the SQS queue. Increase the Lambda function concurrency limit.
- B. Check the `ApproximateAgeOfOldestMessage` metnc for the SQS queue. Configure a redrive policy on the SQS queue.
- C. Check the `NumberOfMessagesSent` metric for the SQS queue. Increase the SQS queue visibility timeout.
- D. Check the `WriteThrottleEvents` metric for the DynamoDB table. Increase the maximum write capacity units (WCUs) for the table's scaling policy.
- E. Check the `Throttles` metric for the Lambda function. Increase the Lambda function timeout.

Suggested Answer: AD

Community vote distribution

AD (100%)

  **madperro** Highly Voted 1 year, 9 months ago

Selected Answer: AD



AD look fine.

upvoted 10 times

  **kiwtirApp** Most Recent 10 months ago

Everyone who has commented AD has not provided any reasoning. Bunch of sheeps in the comments.

upvoted 3 times

  **tg9** 9 months ago

D: This action is important because if the `WriteThrottleEvents` metric is high, it indicates that DynamoDB is throttling writes due to insufficient write capacity. By increasing the maximum WCUs, you ensure that the table can handle the increased write throughput required by the Lambda function, thus reducing delays in order processing.

even though DynamoDB has auto-scaling enabled, it's still important to monitor the `WriteThrottleEvents` metric. Auto-scaling adjusts capacity based on the workload, but it may not always keep up with sudden spikes in demand or be configured optimally for this specific use case.

Ensuring that the maximum write capacity units (WCUs) are set appropriately can help prevent throttling during peak times.

upvoted 4 times



  **thanhv142** 1 year, 1 month ago

A and D is correct:

A: Check `ApproximateAgeOfOldestMessage` and increase concurrency accordingly

D: Check `throttleevent` (the number of rejected requests) and increase max write accordingly.

upvoted 3 times

  **SPRao** 1 year, 3 months ago

Answer should be A & C. D is wrong as DynamoDB has autoscaling enabled so `Writethrottle` should not be the case.

upvoted 3 times

  **harithzainudin** 1 year, 3 months ago

Eventhough the statement say

"The DynamoDB table has auto scaling enabled for read and write capacity."

A and D still are a good answer.

Option C looks at the `NumberOfMessagesSent` metric and suggests increasing the SQS queue visibility timeout. This action is more relevant when messages are not being processed before the visibility timeout expires, but it does not seem to be the primary issue in this scenario.

upvoted 3 times

🗲️ 👤 **Jonfernz** 1 year, 6 months ago

A: Check the ApproximateAgeOfOldestMessage metric for the SQS queue. If this age is high, increasing the Lambda function's concurrency limit would help speed up processing.

D: Check the WriteThrottleEvents metric for the DynamoDB table. If write operations are being throttled, increasing the maximum WCUs for the table's scaling policy could help.

upvoted 4 times

🗲️ 👤 **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: AD

AD are accurate for this scenario.

upvoted 4 times

🗲️ 👤 **alice2020** 1 year, 11 months ago

a and d are correct

upvoted 1 times

🗲️ 👤 **ele** 1 year, 11 months ago

Selected Answer: AD

upscale both

upvoted 2 times

A company has a single AWS account that runs hundreds of Amazon EC2 instances in a single AWS Region. New EC2 instances are launched and terminated each hour in the account. The account also includes existing EC2 instances that have been running for longer than a week. The company's security policy requires all running EC2 instances to use an EC2 instance profile. If an EC2 instance does not have an instance profile attached, the EC2 instance must use a default instance profile that has no IAM permissions assigned.

A DevOps engineer reviews the account and discovers EC2 instances that are running without an instance profile. During the review, the DevOps engineer also observes that new EC2 instances are being launched without an instance profile.


Which solution will ensure that an instance profile is attached to all existing and future EC2 instances in the Region?

- A. Configure an Amazon EventBridge rule that reacts to EC2 RunInstances API calls. Configure the rule to invoke an AWS Lambda function to attach the default instance profile to the EC2 instances.
- B. Configure the ec2-instance-profile-attached AWS Config managed rule with a trigger type of configuration changes. Configure an automatic remediation action that invokes an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
- C. Configure an Amazon EventBridge rule that reacts to EC2 StartInstances API calls. Configure the rule to invoke an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
- D. Configure the iam-role-managed-policy-check AWS Config managed rule with a trigger type of configuration changes. Configure an automatic remediation action that invokes an AWS Lambda function to attach the default instance profile to the EC2 instances.

Suggested Answer: B

Community vote distribution

B (100%)

 **koenigParas2324** Highly Voted 9 months, 3 weeks ago

Selected Answer: B

WS Config, specifically utilizing the "ec2-instance-profile-attached" managed rule with the configuration change trigger type. This rule helps monitor the attachment of instance profiles to EC2 instances. An automatic remediation action can be configured within AWS Config to respond when instances are found without an instance profile attached. The remediation action would execute an AWS Systems Manager Automation runbook to attach the default instance profile to those instances.

upvoted 8 times

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

Selected Answer: B

How is it possible to remember the list of all the config rules?

upvoted 3 times

 **Gomer** 3 months, 2 weeks ago

Selected Answer: B

I concur the best answer seems to be "B". However, I have not been able to trigger exactly what kind of "configuration change" triggers the config rule (e.g. "starting" or "running" an instance isn't a configuration change, but a state change. The real world answer (IMHO) would be to just kick off the AWS Config rule manually or on a schedule. I'd also take steps to ensure that all EC2 launch templates specify an instance profile so I'm not running around trying to fix things that shouldn't have been left broken from the start.

upvoted 1 times

 **Gillar** 5 months ago

Selected Answer: B

The rules AWS Config

upvoted 1 times

 **thanhnv142** 7 months, 2 weeks ago

B is correct: AWS config + runbook is the right way for remediation

upvoted 4 times

 **thanhnv142** 7 months, 1 week ago

B is correct: AWS Config run in combination with SSM Automation run book is the recommended way

A: this option only remediate new instances

C: this option only remediate instances that have been stopped.

D: automatic remediation action should invoke Automation run book, not lambda
upvoted 5 times

🗨️ 👤 **DucSiu** 9 months ago

B

Config: ec2-instance-profile-attached

SSM Automation: AttachedIAMtoinstances

upvoted 3 times

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

Selected Answer: B

B is correct.

<https://docs.aws.amazon.com/config/latest/developerguide/ec2-instance-profile-attached.html>

upvoted 2 times

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

Selected Answer: B

B is correct

upvoted 1 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

correct answer is B

upvoted 1 times

🗨️ 👤 **ele** 1 year, 5 months ago

Selected Answer: B

B , no brainer


upvoted 1 times

A DevOps engineer is building a continuous deployment pipeline for a serverless application that uses AWS Lambda functions. The company wants to reduce the customer impact of an unsuccessful deployment. The company also wants to monitor for issues. Which deploy stage configuration will meet these requirements?

- A. Use an AWS Serverless Application Model (AWS SAM) template to define the serverless application. Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Type. Use Amazon CloudWatch alarms to monitor the health of the functions.
- B. Use AWS CloudFormation to publish a new stack update, and include Amazon CloudWatch alarms on all resources. Set up an AWS CodePipeline approval action for a developer to verify and approve the AWS CloudFormation change set.
- C. Use AWS CloudFormation to publish a new version on every stack update, and include Amazon CloudWatch alarms on all resources. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.
- D. Use AWS CodeBuild to add sample event payloads for testing to the Lambda functions. Publish a new version of the functions, and include Amazon CloudWatch alarms. Update the production alias to point to the new version. Configure rollbacks to occur when an alarm is in the ALARM state.

Suggested Answer: A

Community vote distribution



zolthar_z Highly Voted 1 year, 3 months ago

Certification TIP: 99% of questions regarding lambda and cloudformation the answer is the one that involves SAM
upvoted 21 times

harithzainudin 1 year, 3 months ago

i couldnt agree more with this
upvoted 3 times

Jonfernz Highly Voted 1 year, 6 months ago

Selected Answer: A
A

Reducing Customer Impact: AWS CodeDeploy with Canary deployments (Canary10Percent15Minutes) will incrementally roll out the new version. Initially, 10% of the traffic will be directed to the new version, and if everything goes well, the rest of the traffic will be shifted over the span of 15 minutes. This cautious rollout minimizes the risk and impact on customers.

Monitoring: Amazon CloudWatch alarms can be configured to track function errors, latency, and other important metrics. If anything goes awry, you can act promptly.
upvoted 6 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: A
A

Keywords: Serverless Application related with AWS SAM
upvoted 3 times

zijo 1 year ago

Canary10Percent15Minutes refers to a specific type of deployment strategy used in the context of serverless applications, particularly with tools like AWS SAM (Serverless Application Model).
upvoted 1 times

thanhnv142 1 year, 1 month ago

A is correct: <a continuous deployment pipeline for a serverless application> means AWS SAM.
B, C and D: no mention of SAM
upvoted 4 times

thanhnv142 1 year, 1 month ago

A: use serverless code deployment is the right way
upvoted 1 times

🗳️ 👤 **z_inderjot** 1 year, 2 months ago

Selected Answer: A

A - SAM for Lambda deployment
Reduce Custome Impact - Canary got it covered
upvoted 1 times

🗳️ 👤 **RVivek** 1 year, 5 months ago

Selected Answer: A

A CodeDepoly is for canary deployment , cloudwatch alarm for monitoring, if aram is raised then codedeploy automatically rolls back
D- Using Codebuild for controlled deployment is not good. Codebuild is for build and testing
upvoted 1 times

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

A is the answer: <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>
upvoted 2 times

🗳️ 👤 **Kojhani** 1 year, 7 months ago

Selected Answer: A Canary Deployment
upvoted 3 times

🗳️ 👤 **SanChan** 1 year, 9 months ago

Selected Answer: D

My Answer is D which can help reduce the customer impact of an unsuccessful deployment, also got monitor and rollbacks to occur when an alarm is in the ALARM state.

A, this option does not provide a rollback plan in case of failures, which could further increase the customer impact of a failed deployment.
This strategy can help detect any issues early on, it does not guarantee that the impact on customers will be reduced since some customers might still be affected by the issues.

Someone can tell me more?
upvoted 3 times

🗳️ 👤 **RVivek** 1 year, 6 months ago

Alarms: These are CloudWatch alarms that are triggered by any errors raised by the deployment. When encountered, they automatically roll back your deployment. For example, if the updated code you're deploying causes errors within the application. Another example is if any AWS Lambda or custom CloudWatch metrics that you specified have breached the alarm threshold.

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>
upvoted 1 times

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

codebuild is not use for continous deployment
upvoted 2 times

🗳️ 👤 **SanChan** 1 year, 9 months ago

checked, Answer D does not provide a gradual deployment strategy that reduces the customer impact of a new deployment, which was one of the requirements given in the question.

So the final answer should be A
upvoted 6 times

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

Selected Answer: A

A looks fine.
upvoted 1 times

🗳️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: A



A. Use an AWS Serverless Application Model (AWS SAM) template to define the serverless application. Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Type. Use Amazon CloudWatch alarms to monitor the health of the functions would be the best deploy stage configuration for meeting the requirements of reducing customer impact of an unsuccessful

deployment and monitoring for issues.

Option A uses AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Type, which gradually deploys the new version of the function to a small subset of users before deploying it to the entire fleet. This approach reduces customer impact of an unsuccessful deployment.



Additionally, Amazon CloudWatch alarms are used to monitor the health of the functions, which can provide real-time feedback on any issues that arise. This meets the requirement of monitoring for issues.

upvoted 2 times

  **alce2020** 1 year, 11 months ago

A is the correct answer

upvoted 1 times

  **ele** 1 year, 11 months ago

Selected Answer: A

A looks realistic

upvoted 1 times

To run an application, a DevOps engineer launches an Amazon EC2 instance with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access to the internet. While the instances launch successfully and show as healthy, the application does not seem to be installed.

Which of the following should successfully install the application while complying with the new rule?

- A. Launch the instances in a public subnet with Elastic IP addresses attached. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.
- B. Set up a NAT gateway. Deploy the EC2 instances to a private subnet. Update the private subnet's route table to use the NAT gateway as the default route.
- C. Publish the application artifacts to an Amazon S3 bucket and create a VPC endpoint for S3. Assign an IAM instance profile to the EC2 instances so they can read the application artifacts from the S3 bucket.
- D. Create a security group for the application instances and allow only outbound traffic to the artifact repository. Remove the security group rule once the install is complete.

Suggested Answer: C

Community vote distribution

C (90%) 10%

🗳️ **z_inderjot** Highly Voted 1 year, 2 months ago

Selected Answer: C

C - in the answer

Though we can use both B and C, since we only want to download to package at the time of initialization. So there is no need to have continuous access to internet. Therefore, it is cheap and optimal to use S3.

upvoted 8 times

🗳️ **namtp** Most Recent 7 months, 1 week ago

Selected Answer: C

C is the correct answer.

no access to the internet but connect to aws services => private endpoint

upvoted 1 times

🗳️ **thanhnv142** 1 year, 1 month ago

C is correct: all other options utilize internet connections

upvoted 3 times

🗳️ **harithzainudin** 1 year, 3 months ago

Selected Answer: C

C is the correct one.

all other option will allow internet access which is not compliance with the reqs

upvoted 3 times

🗳️ **zolthar_z** 1 year, 3 months ago

Selected Answer: C

C: Can't be B because with the NAT the EC2 still has internet access

upvoted 3 times

🗳️ **robertohyena** 1 year, 4 months ago

C is the correct answer.

A B D are not correct.

Keywords:

- requires the instances to run with no access to the internet

upvoted 2 times


- 🗳️ 👤 **rowanwally** 1 year, 3 months ago
is the dump still valid?
upvoted 1 times
- 🗳️ 👤 **bosmanx** 1 year, 4 months ago
Selected Answer: C
B is incorrect, the new policy is "no access to the internet"
upvoted 2 times
- 🗳️ 👤 **DevopsNoob** 1 year, 4 months ago
C is the answer. B would enable internet access from the instance.
upvoted 1 times
- 🗳️ 👤 **Ffida** 1 year, 5 months ago
C is correct and B, which is specifically for NAT. in question they have asked that no internet access from the instance, so If we enable NAT then from outside no one can access the instance but internet will be accessible on the instance using NAT.
upvoted 1 times
- 🗳️ 👤 **ataince** 1 year, 5 months ago
C is correct
B: "instances to run with no access to the internet." so you can not use NAT
upvoted 1 times
- 🗳️ 👤 **DaddyDee** 1 year, 6 months ago
C is the answer, you can use artifacts in s3 with vpc endpoints. With a gateway endpoint, you can access Amazon S3 from your VPC, without requiring an internet gateway or NAT device for your VPC, and with no additional cost.
<https://repost.aws/knowledge-center/ec2-systems-manager-vpc-endpoints>
<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html>
upvoted 2 times
- 🗳️ 👤 **rahulsingha2112** 1 year, 6 months ago
C is correct as solution required no internet access
upvoted 1 times
- 🗳️ 👤 **ggrodskiy** 1 year, 7 months ago
Correct C.
upvoted 1 times
- 🗳️ 👤 **madperro** 1 year, 9 months ago
Selected Answer: C
C is the answer. B gives the instances access to the Internet.
upvoted 1 times
- 🗳️ 👤 **rdoty** 1 year, 9 months ago
Selected Answer: C
Def C, all others include access to the internet
upvoted 1 times
- 🗳️ 👤 **ProfXsamson** 1 year, 9 months ago
This is supposed to be a Choose two answer. BC
upvoted 1 times
- 🗳️ 👤 **Akaza** 1 year, 9 months ago
NAT GW for me
upvoted 1 times

A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote main branch as the trigger for the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes. Which of the following actions should be taken to troubleshoot this issue?

- A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.
- B. Check that the CodePipeline service role has permission to access the CodeCommit repository.
- C. Check that the developer's IAM role has permission to push to the CodeCommit repository.
- D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

Suggested Answer: A

Community vote distribution



Dushank Highly Voted 1 year, 5 months ago

Selected Answer: B

A: EventBridge rules are not a requirement for CodePipeline to trigger from a CodeCommit repository. CodePipeline directly integrates with CodeCommit without needing EventBridge.

B: is a likely cause. The CodePipeline service role needs permissions to access the CodeCommit repository in order to start the pipeline execution when new code is pushed.

C: If the developer was able to push code changes to the CodeCommit repository, then their IAM role permissions with respect to CodeCommit are likely fine. This isn't the issue.

D: If the pipeline didn't start, CloudWatch Logs could give insights. However, these logs will only exist if the pipeline actually attempted to start but failed. If the pipeline never started, checking logs won't help.

Given these options, Option B: is the correct answer.

upvoted 18 times

a54b16f Highly Voted 1 year, 1 month ago

Selected Answer: A

B would throw out "Permission denied" error immediately, rather than no reaction for 10 minutes.

upvoted 11 times

spring21 Most Recent 2 months, 2 weeks ago

Selected Answer: A

Change detection:

When you set up a CodePipeline with a CodeCommit source, the default behavior is to use an EventBridge rule to detect changes in the repository, eliminating the need for manual polling mechanisms

upvoted 1 times

ZingjieG87 2 months, 3 weeks ago

Selected Answer: B

The question states that the developer has pushed the change to the repository, which means CodeCommit has no issue, the pipeline doesn't rely on EventBridge to trigger.

upvoted 1 times

Serial_X25 3 months, 3 weeks ago

Selected Answer: A

CodePipeline connects to third-party source providers directly using CodeConnections, <https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-connections.html>, but it should use EventBridge for CodeCommit, <https://docs.aws.amazon.com/codepipeline/latest/userguide/triggering.html>.

upvoted 1 times

🗨️ 👤 **Jonalb** 4 months ago

Selected Answer: A

A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.

This approach directly addresses the most likely cause: a missing or misconfigured EventBridge rule that prevents CodePipeline from starting in response to changes in the CodeCommit repository.

upvoted 1 times

🗨️ 👤 **jamesf** 7 months, 2 weeks ago

Selected Answer: A

A - EventBridge rule is one of the recommended ways to configure CodePipeline to automatically trigger based on changes in a CodeCommit repository

B - if "Permission denied", the error message should prompt immediately, rather than no reaction for 10 minutes for the pipeline. Mean the pipeline not even start

Reference:

<https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-about-starting.html#change-detection-methods>

<https://docs.aws.amazon.com/codepipeline/latest/userguide/tutorials-simple-codecommit.html>

upvoted 3 times

🗨️ 👤 **trungtd** 7 months, 4 weeks ago

Selected Answer: A

EventBridge rule is one of the recommended ways to configure CodePipeline to automatically trigger based on changes in a CodeCommit repository

upvoted 2 times

🗨️ 👤 **Mordans** 8 months, 2 weeks ago

Selected Answer: A

For CodePipeline to be triggered by changes in a CodeCommit repository, an EventBridge rule (formerly CloudWatch Events rule) needs to be set up. This rule listens for specific events (like commits to the main branch) and triggers the pipeline accordingly.

upvoted 1 times

🗨️ 👤 **aefuen1** 8 months, 2 weeks ago

Selected Answer: B

It's B.

upvoted 1 times

🗨️ 👤 **flaacko** 6 months, 3 weeks ago

The answer is not B because if the CodePipeline service linked role didn't have permissions to access CodeCommit, you will get a "Permissions denied" error immediately but the question said you didn't get any reaction in 10 minutes so the only possible scenario we would be dealing with here is not having an EventBridge rule that triggers the pipeline. When you use the console to create or edit a pipeline, the change detection resources are created for you. If you use the AWS CLI to create the pipeline, you must create the additional resources yourself.

Reference: <https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-create.html>

upvoted 1 times

🗨️ 👤 **xdkonorek2** 8 months, 4 weeks ago

Selected Answer: A

the answer is A because if codepipeline has no access to codecommit pipeline is triggered and source stage fails with:

...

The service role or action role doesn't have the permissions required to access the AWS CodeCommit repository named random-repo. Update the IAM role permissions, and then try again

...

upvoted 1 times

🗨️ 👤 **k23319** 9 months, 2 weeks ago

Selected Answer: B

B is right.


upvoted 1 times

🗨️ 👤 **liuyomz** 10 months ago

Selected Answer: B

Just voting to fix the results, because clearly its B, as explained by top 2 comments here.

upvoted 2 times

  **vn_thanhtung** 9 months, 3 weeks ago

B wrong.

upvoted 2 times

  **c3518fc** 10 months ago

Selected Answer: B

The first step in troubleshooting this issue should be to check that the CodePipeline service role has the required permissions to access the CodeCommit repository. If the permissions are correct, then you can proceed with other troubleshooting steps, such as checking the CloudWatch Logs for any errors or failures.

upvoted 1 times

  **c3518fc** 11 months ago

Selected Answer: B

Not sure why most people here are even considering A. CodePipeline does not use Amazon EventBridge to trigger pipeline executions based on changes in CodeCommit repositories. Instead, it directly integrates with CodeCommit and monitors repository changes internally.

upvoted 4 times

  **vn_thanhtung** 10 months, 1 week ago

You are wrong, correct answer is A



upvoted 2 times

  **vn_thanhtung** 10 months, 1 week ago

When you create a pipeline from CodePipeline during the step-by-step it creates a CloudWatch Event rule for a given branch and repo like this:

```
{
  "source": [
    "aws.codecommit"
  ],
  "detail-type": [
    "CodeCommit Repository State Change"
  ],
  "resources": [
    "arn:aws:codecommit:us-east-1:xxxxx:repo-name"
  ],
  "detail": {
    "event": [
      "referenceCreated",
      "referenceUpdated"
    ],
    "referenceType": [
      "branch"
    ],
    "referenceName": [
      "master"
    ]
  }
}
```

upvoted 2 times



  **ogerber** 11 months, 3 weeks ago

Selected Answer: A

"After you select the repository name and branch, a message displays the Amazon CloudWatch Events rule to be created for this pipeline."

<https://docs.aws.amazon.com/codepipeline/latest/userguide/tutorials-simple-codecommit.html>

upvoted 4 times

  **alexleely** 11 months, 4 weeks ago

Selected Answer: A

I highly believe it is A.

Even though CodePipeline directly integrates with CodeCommit, this integration automatically creates a EventBridge rule for you if it is created through the console.

<https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-trigger-source-repo-changes-console.html>

Since we know that there is no reaction from the pipeline, it would mean that it wasn't triggered at all.

B is about permission which would have thrown an error in the console at that stage, but to even start the first stage, it needs to be trigger first which for the case here.

C shouldn't be the answer as the question already said that it was pushed into the repository.

upvoted 3 times

A company's developers use Amazon EC2 instances as remote workstations. The company is concerned that users can create or modify EC2 security groups to allow unrestricted inbound access.

A DevOps engineer needs to develop a solution to detect when users create unrestricted security group rules. The solution must detect changes to security group rules in near real time, remove unrestricted rules, and send email notifications to the security team. The DevOps engineer has created an AWS Lambda function that checks for security group ID from input, removes rules that grant unrestricted access, and sends notifications through Amazon Simple Notification Service (Amazon SNS).

What should the DevOps engineer do next to meet the requirements?

- A. Configure the Lambda function to be invoked by the SNS topic. Create an AWS CloudTrail subscription for the SNS topic. Configure a subscription filter for security group modification events.
- B. Create an Amazon EventBridge scheduled rule to invoke the Lambda function. Define a schedule pattern that runs the Lambda function every hour.
- C. Create an Amazon EventBridge event rule that has the default event bus as the source. Define the rule's event pattern to match EC2 security group creation and modification events. Configure the rule to invoke the Lambda function.
- D. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS services. Configure the Lambda function to be invoked by the custom event bus.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **thanhnv142** Highly Voted 7 months, 2 weeks ago

C is correct:

A: lambda should be invoked by Eventbridge

B: we need to act when there is events, not schedully

D: subscribing to events from ALL AWS services incurs a huge cost

upvoted 6 times

🗳️ 👤 **01037** Most Recent 4 months ago

Selected Answer: C

C of course.

But A seems working, and does Aws Config work in this situation?

upvoted 1 times

🗳️ 👤 **c3518fc** 4 months ago

Selected Answer: C

By creating an EventBridge event rule with the appropriate event pattern and configuring it to invoke the Lambda function, the DevOps engineer can effectively detect security group rule changes in near real-time, remove unrestricted rules, and send notifications to the security team. This solution leverages the event-driven architecture of EventBridge and the serverless execution of AWS Lambda, providing a scalable and efficient way to meet the company's security requirements.

upvoted 2 times

🗳️ 👤 **meriemheni** 8 months, 2 weeks ago

selected answer:C

upvoted 2 times

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

Selected Answer: C

C the default bus includes events from AWS services.

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-event-bus.html>

upvoted 4 times

🗳️ 👤 **bcx** 1 year, 3 months ago

Selected Answer: C

Wrong answers:

A. SNS is used here to send a notification post-facto

- B. The question requires "near real time", an hour is not "near real time"
- D. AWS events come on the default event bus, you do not need a custom event bus
- upvoted 4 times

🗨️ 👤 **Aja1** 1 year, 1 month ago

The default event bus in each account receives events from AWS services.

A custom event bus sends events to or receives events from a different account.

A custom event bus sends events to or receives events from a different Region to aggregate events in a single location.

A partner event bus receives events from a SaaS partner.

upvoted 4 times

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

Selected Answer: C

To meet the requirements, the DevOps engineer should create an Amazon EventBridge event rule that has the default event bus as the source. The rule's event pattern should match EC2 security group creation and modification events, and it should be configured to invoke the Lambda function. This solution will allow for near real-time detection of security group rule changes and will trigger the Lambda function to remove any unrestricted rules and send email notifications to the security team.

upvoted 4 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

C is the answer

upvoted 2 times

🗨️ 👤 **Saga** 1 year, 4 months ago

Selected Answer: C

C. Create an Amazon EventBridge event rule that has the default event bus as the source. Define the rule's event pattern to match EC2 security group creation and modification events. Configure the rule to invoke the Lambda function.

The solution requires near real-time detection of changes to security group rules and immediate action to remove unrestricted rules and send email notifications to the security team. The AWS Lambda function created by the DevOps engineer can perform these actions, but it needs to be invoked whenever a security group rule is modified.

Amazon EventBridge is a serverless event bus service that can receive and process events from various AWS services, including Amazon EC2 and Amazon SNS. An EventBridge event rule with the default event bus as the source can be created to match EC2 security group creation and modification events. This rule can then be configured to invoke the Lambda function, which can remove unrestricted rules and send email notifications to the security team.

upvoted 4 times

🗨️ 👤 **ele** 1 year, 5 months ago

Selected Answer: C

<https://repost.aws/knowledge-center/monitor-security-group-changes-ec2>

upvoted 3 times

A DevOps engineer is creating an AWS CloudFormation template to deploy a web service. The web service will run on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). The DevOps engineer must ensure that the service can accept requests from clients that have IPv6 addresses.

What should the DevOps engineer do with the CloudFormation template so that IPv6 clients can access the web service?

- A. Add an IPv6 CIDR block to the VPC and the private subnet for the EC2 instances. Create route table entries for the IPv6 network, use EC2 instance types that support IPv6, and assign IPv6 addresses to each EC2 instance.
- B. Assign each EC2 instance an IPv6 Elastic IP address. Create a target group, and add the EC2 instances as targets. Create a listener on port 443 of the ALB, and associate the target group with the ALB.
- C. Replace the ALB with a Network Load Balancer (NLB). Add an IPv6 CIDR block to the VPC and subnets for the NLB, and assign the NLB an IPv6 Elastic IP address.
- D. Add an IPv6 CIDR block to the VPC and subnets for the ALB. Create a listener on port 443. and specify the dualstack IP address type on the ALB. Create a target group, and add the EC2 instances as targets. Associate the target group with the ALB.

Suggested Answer: B

Community vote distribution

D (100%)

🗳️ 👤 **levster** Highly Voted 1 year, 9 months ago

D

"To support IPv6, configure your Application Load Balancers or Network Load Balancers with the "dualstack" IP address type. This means that clients can communicate with the load balancers using both IPv4 and IPv6 addresses. In a dual-stack IP address type, the DNS name of the load balancer provides both IPv4 and IPv6 addresses, and creates A and AAAA records respectively. "

<https://docs.aws.amazon.com/whitepapers/latest/ipv6-on-aws/scaling-the-dual-stack-network-design-in-aws.html>

upvoted 6 times

🗳️ 👤 **01037** Most Recent 10 months ago

Selected Answer: D

But why is port 443 necessary?

upvoted 3 times

🗳️ 👤 **flaacko** 6 months, 3 weeks ago

Port 443 is the TCP port for HTTPS which is a secured or encrypted version of HTTP. To enable the ALB handle HTTPS traffic having a listener on port 443 is necessary.

upvoted 2 times

🗳️ 👤 **c3518fc** 10 months ago

Selected Answer: D

The correct answer is D. Add an IPv6 CIDR block to the VPC and subnets for the ALB. Create a listener on port 443. and specify the dualstack IP address type on the ALB. Create a target group, and add the EC2 instances as targets. Associate the target group with the ALB.

upvoted 1 times

🗳️ 👤 **zijo** 1 year ago

Why is the need for port 443 reference on D and D has no reference to private subnet. That makes me think the answer is A, but A has no reference to ALB.

upvoted 2 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

D is correct: use dual stack + listener on 443

A: no mention of the ALB

B: no mention of adding dualstack IP to ALB

C: cannot replace the ALB

upvoted 4 times

🗳️ 👤 **sksegha** 1 year, 2 months ago

definitely D

upvoted 2 times

🗳️ 👤 **Jamshif01** 1 year, 2 months ago

keyword is "Dualstack"

upvoted 2 times

🗳️ 👤 **z_inderjot** 1 year, 2 months ago

Selected Answer: D

D is correct , To enable ALB to deal with Ipv6 requests , vpc should enable for dual stack, by configuring a ipv6 cidr , and ALB subnet should also adhere to the same , by having ipv4 and 6 cidr

B is incorrect , we can assign any public ip to instance , since it is in private subnet .

upvoted 3 times

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

Selected Answer: D

D is the correct answer. C is wrong, we don't need Elastic IPs for a private app.

upvoted 1 times

🗳️ 👤 **Rick365** 1 year, 9 months ago

Selected Answer: D

D is answer

upvoted 1 times

🗳️ 👤 **bcx** 1 year, 9 months ago

I would say it is D. The backend instances serving the data can be IPv4. The ALB should serve IPv6 to the public (which is what is required by the question). So the only place in the VPC that needs IPv6 are the ALB subnets.

upvoted 4 times

🗳️ 👤 **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: D

D is the correct answer in this case

upvoted 1 times

🗳️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: D

To allow IPv6 clients to access the web service running on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB) using an AWS CloudFormation template, the DevOps engineer should choose option D:

Add an IPv6 CIDR block to the VPC and subnets for the ALB. Create a listener on port 443, and specify the dualstack IP address type on the ALB. Create a target group, add the EC2 instances as targets, and associate the target group with the ALB.

The dualstack IP address type enables the ALB to support both IPv4 and IPv6 traffic. By adding an IPv6 CIDR block to the VPC and subnets for the ALB, the VPC automatically assigns an IPv6 address to the ALB.

upvoted 4 times

🗳️ 👤 **gdtypk** 1 year, 10 months ago

D

<https://repost.aws/ja/knowledge-center/elb-configure-with-ipv6>

upvoted 4 times

🗳️ 👤 **alce2020** 1 year, 11 months ago

answer is D

upvoted 1 times

🗳️ 👤 **ele** 1 year, 11 months ago

Selected Answer: D

D right

upvoted 2 times

🗳️ 👤 **lqp0_oqpl** 1 year, 11 months ago

I think D

upvoted 1 times

A company uses AWS Organizations and AWS Control Tower to manage all the company's AWS accounts. The company uses the Enterprise Support plan.

A DevOps engineer is using Account Factory for Terraform (AFT) to provision new accounts. When new accounts are provisioned, the DevOps engineer notices that the support plan for the new accounts is set to the Basic Support plan. The DevOps engineer needs to implement a solution to provision the new accounts with the Enterprise Support plan.

Which solution will meet these requirements?

- A. Use an AWS Config conformance pack to deploy the account-part-of-organizations AWS Config rule and to automatically remediate any noncompliant accounts.
- B. Create an AWS Lambda function to create a ticket for AWS Support to add the account to the Enterprise Support plan. Grant the Lambda function the support:ResolveCase permission.
- C. Add an additional value to the control_tower_parameters input to set the AWSEnterpriseSupport parameter as the organization's management account number.
- D. Set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuration. Redeploy AFT and apply the changes.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ **z_inderjot** Highly Voted 8 months, 3 weeks ago

Selected Answer: D

D check docs

<https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html#enterprise-support-option>

upvoted 11 times

🗳️ **5aga** Highly Voted 1 year, 4 months ago

Selected Answer: D

D. Set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuration. Redeploy AFT and apply the changes.

AWS Organizations is a service that helps to manage multiple AWS accounts. AWS Control Tower is a service that makes it easy to set up and govern secure, compliant multi-account AWS environments. Account Factory for Terraform (AFT) is an AWS Control Tower feature that provisions new accounts using Terraform templates.

To provision new accounts with the Enterprise Support plan, the DevOps engineer can set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuration. This flag enables the Enterprise Support plan for newly provisioned accounts.

upvoted 8 times

🗳️ **thanhnv142** Most Recent 7 months ago

Selected Answer: D

D is correct: < Account Fachttps://www.examtopycs.com/exams/amazon/aws-certified-devops-engineer-professional-dop-c02/view/#tory for Terraform (AFT)> means we need to change AFT config

A: AWS Config conformance pack should be used with SSM automation document to remediate

B and C: irrelevant

upvoted 5 times

🗳️ **madperro** 1 year, 3 months ago

Selected Answer: D

D

<https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html>

upvoted 3 times

🗳️ **haazybanj** 1 year, 4 months ago

Selected Answer: D

D. Set the `aft_feature_enterprise_support` feature flag to `True` in the AFT deployment input configuration, and then redeploy AFT to apply the changes. This flag is used to enable the Enterprise Support plan for new accounts provisioned by AFT. By default, AFT provisions accounts with the Basic Support plan. Therefore, enabling this flag will provision accounts with the Enterprise Support plan.

upvoted 3 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

D it is

upvoted 2 times

🗨️ 👤 **ele** 1 year, 5 months ago

Selected Answer: D

D: To enable the Enterprise Support option, set the following feature flag to `True` in your AFT deployment input configuration.

`aft_feature_enterprise_support=true`

<https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html>

upvoted 1 times

🗨️ 👤 **lqp0_0qpl** 1 year, 5 months ago

Why not C?

upvoted 1 times

🗨️ 👤 **bugincloud** 11 months, 3 weeks ago

check this out https://controltower.aws-management.tools/automation/aft_setup/


upvoted 1 times

A company's DevOps engineer uses AWS Systems Manager to perform maintenance tasks during maintenance windows. The company has a few Amazon EC2 instances that require a restart after notifications from AWS Health. The DevOps engineer needs to implement an automated solution to remediate these notifications. The DevOps engineer creates an Amazon EventBridge rule. How should the DevOps engineer configure the EventBridge rule to meet these requirements?

- A. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance. Target a Systems Manager document to restart the EC2 instance.
- B. Configure an event source of Systems Manager and an event type that indicates a maintenance window. Target a Systems Manager document to restart the EC2 instance.
- C. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.
- D. Configure an event source of EC2 and an event type that indicates instance maintenance. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

Suggested Answer: A

Community vote distribution



MarDog Highly Voted 1 year, 8 months ago

And AWS Training and Certification has A as the correct answer in the practice exam.
upvoted 30 times

Seoyong Highly Voted 1 year, 6 months ago

Selected Answer: A

It doesn't need to invoke Lambda.
There is a SSM document , RestartEC2Instance
<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html>
upvoted 7 times

VerRi Most Recent 4 months ago

Selected Answer: A

Bad wording. The SSM document here means the automation document (Runbook), not the Command document. EventBridge + SSM Automation (automation document aka Runbook) is a good practice
upvoted 1 times

jamesf 7 months, 2 weeks ago

Selected Answer: A

A
Using SSM document to restart EC2 Instance. Not require to invoke Lambda.

<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html#automating-instance-actions>
upvoted 1 times

trungtd 8 months, 1 week ago

Selected Answer: A

No need to invoke Lambda.
upvoted 1 times

xdkonorek2 8 months, 2 weeks ago

Selected Answer: C

I'm hesitant between A and C but I'm voting C
1) SSM document is not a valid target, valid targets for SSM are: Automation, Run Command, Opsitem
2) If company is already using maintenance windows devops engineer should use them instead of restarting instances immediately
upvoted 1 times

flaacko 6 months, 3 weeks ago

From the question, the company is already using SSM so there is really no need to create a custom Lambda function. SSM is a valid target action for EventBridge events. You can trigger the running of the AWS-RestartEC2Instance automation document with an EventBridge event which means SSM documents are a valid target.

upvoted 1 times

🗳️ 👤 **4bc91ae** 9 months, 2 weeks ago

Selected Answer: A

easiest way to do this

upvoted 1 times

🗳️ 👤 **Gomer** 9 months, 2 weeks ago

Selected Answer: A

SSM Runbook: AWS-RestartEC2Instance (restart one or more EC2 instances)

upvoted 1 times

🗳️ 👤 **Gomer** 9 months, 2 weeks ago

In reading through some of the responses I think "maintenance windows" (plural) doesn't imply scheduling through Lambda. A DevOps engineer can disable automation during production hours. The scenario is unclear if they want this running all the time, or just enabled to run ONLY in a maintenance window. What I'm sure of is they are wanting the SSM runbook as the answer. In the real world, if productin EC2 instance has a health issue, you might just very well want to reboot it automatically if that truly fixes the problem. Nuff said.

upvoted 1 times

🗳️ 👤 **bont** 9 months, 3 weeks ago

The answer is C because A. This option is incorrect because AWS Health notifications do not trigger Systems Manager maintenance windows directly. Additionally, Systems Manager documents cannot restart EC2 instances directly; they need to be executed through other services like Systems Manager Automation or AWS Lambda.

upvoted 1 times

🗳️ 👤 **zijo** 1 year ago

Answer is A. You can create a maintenance window in AWS SSM and associate the EventBridge rule with the maintenance window. No need to customize the solution with lambda.

upvoted 3 times

🗳️ 👤 **vn_thanhtung** 10 months, 3 weeks ago

A say Target a"Systems Manager document" not support by EB => need to use Lambda => Answer is C

<https://docs.aws.amazon.com/eventbridge/latest/userguide/eb-targets.html>

upvoted 1 times

🗳️ 👤 **vn_thanhtung** 9 months, 3 weeks ago

A Systems Manager document defines the actions that Systems Manager performs on your managed instances. An automation document is a type of Systems Manager document that's used to perform common maintenance and deployment tasks. This includes creating or updating an Amazon Machine Image (AMI). This topic outlines how to create, edit, publish, and delete automation documents with AWS Toolkit.

sorry my mistake, ans is A

upvoted 1 times

🗳️ 👤 **kyuhuck** 1 year ago

Selected Answer: C

Thus, Option C is the most accurate and effective solution for automating EC2 instance restarts in response to AWS Health notifications, leveraging the combined capabilities of AWS Health, Amazon EventBridge, AWS Lambda, and AWS Systems Manager.

upvoted 1 times

🗳️ 👤 **01037** 10 months ago

Why does Lambda have to be involved?

upvoted 2 times

🗳️ 👤 **vortegon** 1 year, 1 month ago

Selected Answer: A

<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html#automating-instance-actions>

upvoted 4 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

A is correct:

B: AWS health should be the event source, not system manager

C and D: should not use lambda if already have System manager

upvoted 2 times

🗨️ 👤 **a54b16f** 1 year, 1 month ago

The system is already using SSM to manage EC2 instances, why would you create another solution and use Lambda ? The maintenance window is added to confuse people. The event is from AWS health and need attention immediately. option A fits perfectly.

upvoted 3 times

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

<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html#automating-instance-actions>

upvoted 2 times

🗨️ 👤 **koenigParas2324** 1 year, 3 months ago

Selected Answer: A

Option A appears to be the most suitable:

Configuring AWS Health as the event source ensures notifications related to EC2 instances are captured.

Targeting a Systems Manager document to restart the EC2 instance aligns with Systems Manager's capabilities for automated tasks like instance restarts.

Option B focuses on Systems Manager events related to maintenance windows, which might not directly align with notifications triggered by AWS Health for EC2 instance maintenance.

upvoted 3 times

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

Selected Answer: A

Answer is A, lets breakdown the question. The first part is the DevOps uses system manager for maintenance windows (ok, normal approach) Second part of the question, some EC2 instances requires a restart after AWS Health notification (So, If there is a AWS Health notification the EC2 instance needs a restart), third part of the question, the DevOps should solve the part 2 problem automatically (but it doesn't say when, only a restart is needed), so .. the first part of the question is a catfish, you need to solve the problem automatically an the best way to do it is the A option,

upvoted 4 times

A company has containerized all of its in-house quality control applications. The company is running Jenkins on Amazon EC2 instances, which require patching and upgrading. The compliance officer has requested a DevOps engineer begin encrypting build artifacts since they contain company intellectual property.

What should the DevOps engineer do to accomplish this in the MOST maintainable manner?

- A. Automate patching and upgrading using AWS Systems Manager on EC2 instances and encrypt Amazon EBS volumes by default.
- B. Deploy Jenkins to an Amazon ECS cluster and copy build artifacts to an Amazon S3 bucket with default encryption enabled.
- C. Leverage AWS CodePipeline with a build action and encrypt the artifacts using AWS Secrets Manager.
- D. Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on EC2 instances.

Suggested Answer: D

Community vote distribution

D (84%)

B (16%)

 **sb333** Highly Voted 1 year, 8 months ago

Selected Answer: D


The question wants you to know which solution is the easiest to maintain. It's important not to get thrown by information provided about their current environment. Only the question they ask matters. The question asks which solution is the easiest to "maintain". The question did not ask whether it would be easy to transition from one solution to another or ask you to leverage containers like other parts of their environment.

As a managed service, AWS CodeBuild does not require patching and upgrading. AWS CodeBuild, using Amazon S3, provides automatic artifact encryption. So this solution is the easiest to maintain of all the solutions listed.

<https://docs.aws.amazon.com/codebuild/latest/userguide/welcome.html>

<https://docs.aws.amazon.com/codebuild/latest/userguide/security-encryption.html>


upvoted 22 times

 **madperro** Highly Voted 1 year, 9 months ago

Selected Answer: D

While B will require less changes to the build process I assume AWS is promoting managed services here and expects D answer.

upvoted 11 times

 **Ravi_Bulusu** Most Recent 3 months, 3 weeks ago

The answer is B

Containerized Jenkins on ECS:

By deploying Jenkins on Amazon ECS (Elastic Container Service), you can leverage containerized environments to easily scale and manage Jenkins. This reduces the operational overhead of patching and upgrading EC2 instances running Jenkins.

Artifact Storage with Encryption: Storing build artifacts in Amazon S3 with default encryption enabled ensures that all files in the bucket are automatically encrypted at rest using either SSE-S3 or SSE-KMS. This complies with the requirement to protect intellectual property by ensuring encryption of artifacts.


This approach ensures a fully managed and scalable solution for both Jenkins (containerized) and the artifact storage, aligning with best practices for security and compliance.

upvoted 1 times

 **newpotato** 5 months, 2 weeks ago

while option D could be easier for simple projects or when starting from scratch, it may not be the most maintainable solution for a company that already has a significant investment in Jenkins. Option B provides a balanced approach, leveraging Jenkins' capabilities while improving infrastructure management and security.

upvoted 1 times

 **HarryLy** 9 months, 1 week ago

Selected Answer: D

AWS codebuild use kms encryption key by default

upvoted 1 times

 **Gomer** 9 months, 2 weeks ago

Selected Answer: D

"D" for me based on sb333's comments, etc.

upvoted 1 times

🗨️ 👤 **01037** 9 months, 4 weeks ago

Selected Answer: D

D isn't cost effective, but most maintainable

upvoted 1 times

🗨️ 👤 **zijo** 1 year ago

Answer is D

AWS CodeBuild can be seamlessly integrated with containerized applications deployed on Amazon ECS.

AWS CodeBuild utilizes multiple layers of encryption to safeguard your data at rest, in transit, and during execution.

upvoted 1 times

🗨️ 👤 **Vitalydt** 1 year ago

Selected Answer: D

D Seems the best option

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

D is correct: codebuild has encryption by default -> easiest to maintain

A: No mention of encrypting build artifacts

B: Amazon S3 exryption only protect data at rest, not encrypting the data

C: Using both AWS codepipeline and AWS secret manager incurs more costs and makes maintenance much more difficult

upvoted 3 times

🗨️ 👤 **DucSiu** 1 year, 2 months ago

D is the right answer

upvoted 1 times

🗨️ 👤 **Sazeka** 1 year, 3 months ago

Selected Answer: D

D is the right answer

upvoted 1 times

🗨️ 👤 **2pk** 1 year, 4 months ago

Selected Answer: B

B is the answer . The ask is not to re engineer the whole solution it's just a simple task which needs encrypt the artifact.

Jenkins on Amazon ECS: Running Jenkins in an Amazon ECS cluster allows you to containerize your Jenkins setup, making it easier to manage and scale. ECS offers high availability, scalability, and easy maintenance.

Normally Jenkin should run on ECS so it can handle multiple agents while use S3 as the default encryption.

upvoted 1 times

🗨️ 👤 **RVivek** 1 year, 5 months ago

Selected Answer: D

MOST maintainable manner is repacing jenkins with Codebuild a fully managed service

If the question had been with minimal chnage to the envornment then B would be best

upvoted 2 times

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

Answer is D: MOST maintainable manner/managed service is the key word and there is no need to patch and upgrade. There is ECS with EC2 instances and ECS with fargate and the question is not explicit. Hence maintenance wise, a managed service is the way to go.

<https://jenkinshero.com/jenkins-vs-aws-codebuild-for-building-docker-images/>

upvoted 1 times

🗨️ 👤 **habros** 1 year, 8 months ago

Selected Answer: D

Technically CodeBuild runs on a VM... albeit disposable. Switching on EC2 24/7 is not cost effective either.

upvoted 1 times

🗨️ 👤 **tartarus23** 1 year, 8 months ago

Selected Answer: D

D. Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on EC2 instances.

Explanation: AWS CodeBuild is a fully managed continuous integration service that compiles source code, runs tests, and produces software packages that are ready to deploy. With CodeBuild, you don't need to provision, manage, and scale your own build servers. It also provides built-in support for artifact encryption, which would satisfy the compliance officer's requirements. This would eliminate the need for patching and upgrading Jenkins on EC2 instances, as well as the need to handle encryption at the storage level.

upvoted 2 times

An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.

All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.

How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

- A. Add a DeletionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
- B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM role. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.
- C. Identify the resource that was not deleted. Manually empty the S3 bucket and then delete it.
- D. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resource. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **thanhnv142** Highly Voted 7 months, 2 weeks ago

B is correct:

- Cant delete S3 so must check S3
 - There are several DeletionPolicy option in ACF: delete, retain, snapshot. For S3, even if there is delete flag, S3 can only be deleted if all objects are removed
 - A: wrong - add delete flag to deletionpolicy cant forcing deletion of S3
 - C: should not manually do the task
 - D: should not swap to AWS opsworks
- upvoted 7 times

🗳️ 👤 **n_d1** Highly Voted 1 year, 2 months ago

B. As per the AWS DeletionPolicy Options documentation it says, "For Amazon S3 buckets, you must delete all objects in the bucket for deletion to succeed."

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

upvoted 7 times

🗳️ 👤 **HarryLy** Most Recent 3 months ago

Selected Answer: B

Cloudformation does not have any behavior to force delete not empty bucket, need to invoke a custom lambda function to delete it

upvoted 1 times

🗳️ 👤 **c3518fc** 4 months ago

Selected Answer: B

Keyword "Custom Resource"

upvoted 1 times

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

Selected Answer: B

B is a correct answer. A is wrong, you can't delete a bucket that has any objects.

upvoted 1 times

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

Selected Answer: B

B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM role. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.

upvoted 2 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

B is the correct answer

upvoted 1 times

🗨️ 👤 **ele** 1 year, 5 months ago

Selected Answer: B

Because it's B. CFN will not delete non-empty bucket. It must be emptied first. Custom resource will do it.

upvoted 3 times

🗨️ 👤 **lqp0_oqpl** 1 year, 5 months ago

Why not A?

upvoted 1 times

🗨️ 👤 **tycho** 1 year, 4 months ago

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

deletion policy seems fine as well ...

upvoted 1 times

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

As per the linked article: "For Amazon S3 buckets, you must delete all objects in the bucket for deletion to succeed."

upvoted 2 times

A company has an AWS CodePipeline pipeline that is configured with an Amazon S3 bucket in the eu-west-1 Region. The pipeline deploys an AWS Lambda application to the same Region. The pipeline consists of an AWS CodeBuild project build action and an AWS CloudFormation deploy action.

The CodeBuild project uses the `aws cloudformation package` AWS CLI command to build an artifact that contains the Lambda function code's .zip file and the CloudFormation template. The CloudFormation deploy action references the CloudFormation template from the output artifact of the CodeBuild project's build action.


The company wants to also deploy the Lambda application to the us-east-1 Region by using the pipeline in eu-west-1. A DevOps engineer has already updated the CodeBuild project to use the `aws cloudformation package` command to produce an additional output artifact for us-east-1.

Which combination of additional steps should the DevOps engineer take to meet these requirements? (Choose two.)

- A. Modify the CloudFormation template to include a parameter for the Lambda function code's zip file location. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to pass in the us-east-1 artifact location as a parameter override.
- B. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.
- C. Create an S3 bucket in us-east-1. Configure the S3 bucket policy to allow CodePipeline to have read and write access.
- D. Create an S3 bucket in us-east-1. Configure S3 Cross-Region Replication (CRR) from the S3 bucket in eu-west-1 to the S3 bucket in us-east-1.
- E. Modify the pipeline to include the S3 bucket for us-east-1 as an artifact store. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.

Suggested Answer: AB

Community vote distribution



madperro Highly Voted 1 year, 9 months ago

Selected Answer: CE

As below. You need S# bucket in the new region so C. You need to output artifacts to this new bucket so E.

<https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-create-cross-region.html>

upvoted 19 times

steli0 Most Recent 3 months, 2 weeks ago

Selected Answer: CE

Artifact bucket is needed in the deployment region

upvoted 1 times

Ravi_Bulusu 3 months, 3 weeks ago

The Answer is : AB

To deploy to a different region, the CloudFormation template should be flexible enough to accept a parameter for the Lambda function's zip file location. This allows the template to be reused in both regions. The new CloudFormation action in us-east-1 should reference this parameter and pass in the appropriate location of the artifact for that region

After the CodeBuild project outputs artifacts for both eu-west-1 and us-east-1, you need a separate CloudFormation deploy action in the pipeline that targets the us-east-1 region. This action should reference the CloudFormation template from the us-east-1 output artifact produced by the CodeBuild step.

upvoted 1 times

jamesf 7 months, 2 weeks ago

Selected Answer: CE

CE

Not D because "A DevOps engineer has already updated the CodeBuild project to use the `aws cloudformation package` command to produce an additional output artifact for us-east-1"

<https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-create-cross-region.html>

upvoted 1 times

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

C and E are the right answers.

upvoted 1 times

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

A: It suggests pointing directly to the Lambda function but we also need the Cloudformation template. So we can rule option A out.

B: Here, we are missing the new region (us-east-1) artifact store where the new Cloudformation deploys action would store the artifacts upon completion.

E: Includes the missing part in option B and is the right answer.

D: "A DevOps engineer has already updated the CodeBuild project to use the AWS CloudFormation package command to produce an additional output artifact for us-east-1." So as we are directly producing the artifact in the S3 bucket in us-east-1, I don't see the point of having a cross-replication.

C: This option is mandatory as we should provide permissions to services (CodePipeline) to access resources (S3 bucket).

upvoted 1 times

🗳️ 👤 **ihustle** 9 months, 2 weeks ago

B and C are the answers.

The two important things to note here are the use of AWS CLI and artifacts from two different regions.

upvoted 1 times

🗳️ 👤 **ihustle** 9 months, 2 weeks ago

Apologies, I meant C and E.

upvoted 1 times

🗳️ 👤 **kyuhuck** 1 year ago

Selected Answer: AB

a/b correct

a: the cloudformation template should be modified to include a parameter that indicates the location of the .zip file containing the lambda function's code, this allows the cloudformation deploy action to use the correct artifact depending on the region, this is critical because lambda functions need to reference their code artifacts from the same region they are being deployed in, b. you would also need to create a new cloudformation deploy action from the us-east-1 region within the pipeline this action should be configured to use the cloudformation template from the artifact that was specifically created for us-east-1

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

C and E are correct: To achieve the goal, we need an empty S3 in the us-east-1 and create additional stage in the pipeline

A: No mention of S3 - incorrect

B: no mention of S3 - incorrect

D: we need an empty S3 to store artifact, Cross-Region Replicate incurs more unnecessary cost. Additionally, this way forces the S3 in us-east-1 to be exactly like that of us-west-1, which is incorrect. Each S3 has a different set of artifacts (though they might be very similar)

upvoted 4 times

🗳️ 👤 **Bans** 1 year, 2 months ago

Why C and not D?

upvoted 1 times

🗳️ 👤 **tgx** 7 months, 4 weeks ago

"A DevOps engineer has already updated the CodeBuild project to use the aws cloudformation package command to produce an additional output artifact for us-east-1"

upvoted 2 times

🗳️ 👤 **DucSiu** 1 year, 2 months ago

<https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-create-cross-region.html#actions-create-cross-region-cfn>

CE is my answer

upvoted 1 times

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

For CloudFormation you need to add the Region parameter: <https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-create-cross-region.html#actions-create-cross-region-cfn>

upvoted 1 times

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

Selected Answer: CE

Answers: C E

Scenario:

- We have Pipeline in RegionA (eu-west-1 Region)
- We have Deploy action in RegionA (eu-west-1 Region)

Requirements:

- Need to have Deploy to RegionB (us-east-1 Region)
- And still use RegionA pipeline above (eu-west-1 Region)

upvoted 3 times

  **robertohyena** 1 year, 3 months ago

Which combination of additional steps to meet requirements:

- Create bucket in RegionB (us-east-1 Region) [artifact store] This will be the OutputArtifact bucket for Deploy action in RegionB (us-east-1 Region)
- Add a cross-Region action to a pipeline (CLI)
- Described in step #3 "Modify the pipeline to include the S3 bucket for us-east-1 as an artifact store."
- Described in step #2 "Create a new CloudFormation deploy action for us-east-1 in the pipeline." and "Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact."

REF: <https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-create-cross-region.html#actions-cross-region-cli>

INCORRECT

- A B cannot be a "combination of steps". They both have "Create a new CloudFormation deploy action for us-east-1 in the pipeline."
- D - we do not need S3 Cross-Region Replication (CRR) in the solution.

upvoted 2 times

  **2pk** 1 year, 4 months ago

A and B is the answer. Anything related to create S3 is wrong since Code Deploy have ability to automatically share artifacts to other regions

upvoted 1 times

  **z_inderjot** 1 year, 2 months ago

we are not using codedeploy

upvoted 3 times

  **zain1258** 1 year, 4 months ago

Selected Answer: CE



C & E are correct options

upvoted 3 times

  **dencccc** 1 year, 4 months ago

Would go for CE

upvoted 1 times

  **DZ_Ben** 1 year, 4 months ago

It should be BC! In order to do cross-region deployment, we should create two s3 for each region for ArtifactStore. Then CodeBuild should bundle the sam template and upload to each bucket. Finally CodePipeline should have two separate action for Cloudformation deployment, and each deployment has a region attribute to be defined and needs to pull the template from the ArtifactStore respectively.

upvoted 1 times



  **kacsabacsi78** 1 year, 4 months ago

Selected Answer: AB

C, D and E answers are wrong. CodePipeline automatically creates an S3 bucket in the cross-region for the artifacts. CodePipeline handles the copying of artifacts from one AWS Region to the other Regions when performing cross-region actions.

<https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-create-cross-region.html>

upvoted 2 times

  **Cappy46789** 9 months, 3 weeks ago

That link also says if you are using Cloudformation or CLI then you have to provide the buckets. So C and E


upvoted 1 times

A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive. Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metric. Use the recover action to stop and start the instance. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- B. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instance. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
- C. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failure. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- D. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resource. Add a command in UserData to retrieve the application metadata from Amazon S3.

Suggested Answer: B

Community vote distribution



Jonfernz Highly Voted 1 year, 6 months ago

Selected Answer: B

Both Amazon CloudWatch's recover action and EC2 Auto Recovery are designed to respond to system status check failures, not instance status check failures. System status check failures indicate issues with the underlying hardware, while instance status check failures are often related to issues within your instance (like an OS-level issue).

If the requirement is to handle unresponsiveness due to both system-level and instance-level issues, neither option A nor C would fully meet the requirement. In that case, AWS OpsWorks with auto healing (Option B) could be a better fit since OpsWorks allows you to configure more complex health checks and could recover from both system-level and instance-level issues.

So, if you want to handle both types of unresponsiveness, Option B would be the most comprehensive solution.

upvoted 12 times

flaacko 6 months, 3 weeks ago

May I add that AWS Opswork offers lifecycle events which you can leverage to execute custom actions on the EC2 instance for example retrieving metadata from S3 as the question requested.

upvoted 2 times

endian675 Most Recent 3 months ago

Selected Answer: B

OpsWorks has now been retired, so don't expect to see this question. However, the answer appears to be B.

A: doesn't make sense because S3 notifications only happen if the S3 objects are modified.

C: same argument as A

D: impossible.

upvoted 2 times

BrusingWayne 3 months, 3 weeks ago

Every options are wrong at this moment. Opsworks reached EOL. Other options do not make any sense.

upvoted 1 times

Ravi_Bulusu 3 months, 3 weeks ago

The best approach is C, using EC2 Auto Recovery to monitor and recover the instance if it becomes unresponsive, combined with S3 event notifications to ensure the application metadata is properly retrieved after the instance is back online.

upvoted 1 times

HarryLy 9 months, 1 week ago

Selected Answer: B

B seem correct

upvoted 1 times

🗨️ **Gomer** 9 months, 2 weeks ago

Identical with Question #: 102

upvoted 1 times

🗨️ **hoazgagh** 11 months ago

Selected Answer: B

To automatic restart, must pull artifact for proactive

upvoted 1 times

🗨️ **thanhv142** 1 year, 1 month ago

B: is correct: AWS opworks auto healing will monitor the healthiness of EC2. If there is failure, restart EC2 and pull data from S3 to EC2

A: incorrect because no mention of method to trigger S3 and S3 will not trigger by itself

C: incorrect because no mention of method to trigger S3 and S3 will not trigger by itself

D: Cloud formation only for deploy, this task is about opwork

upvoted 3 times

🗨️ **z_inderjot** 1 year, 2 months ago

Selected Answer: B

OpWorks is deprecated now , So will it be part of exam ? What is the point of learning of service that are not , going to use.

upvoted 3 times

🗨️ **TheAWSRhino** 1 year, 4 months ago

Selected Answer: B

OpWorks is EOL now, however, I think this is the correct answer currently.

upvoted 3 times

🗨️ **harithzainudin** 1 year, 3 months ago

yes indeed. its EOL

upvoted 1 times

🗨️ **beanxyz** 1 year, 6 months ago

Selected Answer: B

A and C are wrong because S3 event notification destination is lambda, sqs and sns topic, you can't directly push metadata to EC2;

D is wrong because although user data can retrieve s3 metadata, it can't restart automatically.

upvoted 4 times

🗨️ **n_d1** 1 year, 8 months ago

Selected Answer: B

B. It doesn't make sense for an S3 event notification to be triggered by an EC2 instance being restarted. The OpsWorks autohealing capability can detect failed instances and replace them.

After the auto-healed instance is back online, OpsWorks triggers a Configure lifecycle event on the instance. The metadata from S3 could be retrieved by the lifecycle event with a recipe.

<https://docs.aws.amazon.com/opsworks/latest/userguide/workinginstances-autohealing.html>

<https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html>

https://github.com/awsdocs/aws-opsworks-user-guide/blob/master/doc_source/create-custom-configure.md

upvoted 2 times

🗨️ **madperro** 1 year, 9 months ago

Selected Answer: B

B, not simplest one but the only that meets requirements.

For A - how can you push data from S3 to EC2? Data needs to be pulled from EC2.

upvoted 2 times

🗨️ **Akaza** 1 year, 9 months ago

Selected Answer: A

A for me

By creating a CloudWatch alarm for the StatusCheckFailed metric, the system can detect if the instance becomes unresponsive. The recover action can then be triggered to automatically stop and start the instance, ensuring it restarts or relaunches when necessary.

Additionally, an S3 event notification can be set up to push the metadata to the instance once it is back up and running. This ensures that the application metadata is retrieved and available after the restart

upvoted 1 times

🗨️ 👤 **bcx** 1 year, 9 months ago

The second part would not work. An S3 notification event occurs only when actions occur on the object. When you restart the instance, nobody is overwriting the object to trigger the notification. IMHO.

upvoted 3 times

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

Selected Answer: B

B seems to be more feasible in this case.

upvoted 2 times

🗨️ 👤 **Mail1964** 1 year, 10 months ago

<https://aws.amazon.com/about-aws/whats-new/2022/03/amazon-ec2-default-automatic-recovery/>

upvoted 1 times

🗨️ 👤 **alce2020** 1 year, 11 months ago

I'd say the answer is A ..you can configure Amazon CloudWatch to monitor the EC2 instance and trigger an automatic restart or relaunch if it becomes unresponsive. You can set up a CloudWatch alarm to monitor the instance's CPU utilization, network traffic, or other metrics, and define an action to take if the alarm is triggered, such as rebooting the instance or terminating and relaunching it.

upvoted 1 times

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

"Use an S3 event notification to push the metadata to the instance when the instance is back up and running." makes no sense

upvoted 1 times

A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.

The attribute mapping list contains two entries. The department key is mapped to `${path:enterprise.department}`. The costCenter key is mapped to `${path:enterprise.costCenter}`.

All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.

Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?

A.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "aws:TagKeys": ["department"]
  }
}
```

B.

```
"Condition": {
  "StringEquals": {
    "aws:PrincipalTag/department": "${aws:ResourceTag/department}"
  }
}
```

C.

```
"Condition": {
  "StringEquals": {
    "ec2:ResourceTag/department": "${aws:PrincipalTag/department}"
  }
}
```

D.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "ec2:ResourceTag/department": ["d1", "d2", "d3"]
  }
}
```

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ **thanhnv142** Highly Voted 1 year, 1 month ago

C is correct: check the EC2's department tag, if it is the same as user(principaltag)'s department tag, allow access.

A: wrong syntax, should be StringEquals only

B: we checking the tag of Ec2, not aws.

D: if config like this, every cases will match and everyone can access every EC2, regardless of department

upvoted 6 times

🗉 👤 **jamesf** Most Recent 7 months, 2 weeks ago

Selected Answer: C

C, related with ABAC.

upvoted 1 times

🗉 👤 **madperro** 1 year, 9 months ago

Selected Answer: C

C, see an example at

<https://docs.aws.amazon.com/singlesignon/latest/userguide/configure-abac.html>

upvoted 4 times

🗉 👤 **alce2020** 1 year, 11 months ago

C is the correct answer

upvoted 2 times

🗉 👤 **ele** 1 year, 11 months ago

Selected Answer: C

<https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/>

upvoted 4 times


A company hosts a security auditing application in an AWS account. The auditing application uses an IAM role to access other AWS accounts. All the accounts are in the same organization in AWS Organizations.

A recent security audit revealed that users in the audited AWS accounts could modify or delete the auditing application's IAM role. The company needs to prevent any modification to the auditing application's IAM role by any entity other than a trusted administrator IAM role. Which solution will meet these requirements?

- A. Create an SCP that includes a Deny statement for changes to the auditing application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the SCP to the root of the organization.
- B. Create an SCP that includes an Allow statement for changes to the auditing application's IAM role by the trusted administrator IAM role. Include a Deny statement for changes by all other IAM principals. Attach the SCP to the IAM service in each AWS account where the auditing application has an IAM role.
- C. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the permissions boundary to the audited AWS accounts.
- D. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the permissions boundary to the auditing application's IAM role in the AWS accounts.

Suggested Answer: C

Community vote distribution



jqso234 1 year, 4 months ago

Selected Answer: A

SCPs (Service Control Policies) are the best way to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

upvoted 20 times

4555894 6 months ago

Selected Answer: A

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html?icmpid=docs_orgs_console

upvoted 1 times

zijo 6 months, 2 weeks ago

Service Control Policies (SCPs) in AWS Organizations can be used to enforce maximum permissions for member accounts. They don't directly grant permissions or create permission boundaries. So C & D can be ruled out.

upvoted 1 times

dzn 6 months, 3 weeks ago

Selected Answer: A

SCPs are applied at the account or OU level and affect all IAM entities within that organization. IAM Permission boundaries are applied individually to specific IAM roles or users.

upvoted 3 times

thanhnv142 7 months, 1 week ago

Selected Answer: A

A is correct: < prevent any modification to the auditing application's IAM role> means scp

A: <Include a condition that allows the trusted administrator IAM role> this is not the same as allow statement. So this option still valid

B: SCP does not have allow statement

C and D: These options make modification to permission boundary of the auditing application's IAM role, which is irrelevant. Other accounts may or may not assume this role.

upvoted 1 times

thanhnv142 7 months, 1 week ago



B is not correct because can only attach scp to AWS org

upvoted 1 times

  **vn_thanhtung** 3 months, 4 weeks ago

B wrong because SCP not support principals

upvoted 1 times

  **flameme** 11 months, 2 weeks ago

AWS supports permissions boundaries for IAM entities (users or roles)

upvoted 1 times

  **aussiehoa** 1 year, 1 month ago

in option A, shouldn't the first half override the second half. Explicitly deny everybody(will not matter if later it says Allow Admin).

upvoted 2 times

  **nlw** 10 months, 1 week ago

I think its because its not two policies. Its only one policy which applies when condition is account not equal security admin account. So A should work

upvoted 4 times

  **ogwu2000** 1 year, 1 month ago

A seems ok. For C its wrong as you don't use permission boundary to deny permission. You use it to specify what and what can be done and not what cannot be done.


upvoted 1 times

  **madperro** 1 year, 3 months ago

Selected Answer: A

For AWS Organizations the SCP is the way to go. So A.



upvoted 1 times

  **bcx** 1 year, 3 months ago

Selected Answer: A

An SCP would accomplish efficiently the task for all the accounts from a single place. A permission boundary is not for that, it would have to be configured in each account and for all the users IMHO.

upvoted 2 times

  **rdoty** 1 year, 3 months ago

Selected Answer: A

It is A without a question. SCP is far more efficient.

upvoted 1 times

  **qan1257** 1 year, 3 months ago

Selected Answer: A

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html?icmpid=docs_orgs_console



upvoted 1 times

  **ParagSanyashiv** 1 year, 4 months ago

Selected Answer: C

C is more suitable option here to restrict permission.

upvoted 1 times


  **5aga** 1 year, 4 months ago

Selected Answer: C

A permissions boundary is designed to restrict permissions on IAM principals, such as roles, such that permissions don't exceed what was originally intended. The permissions boundary uses an AWS or customer managed policy to restrict access, and it's similar to other IAM policies you're familiar with because it has resource, action, and effect statements. A permissions boundary alone doesn't grant access to anything. Rather, it enforces a boundary that can't be exceeded, even if broader permissions are granted by some other policy attached to the role.

<https://aws.amazon.com/blogs/security/when-and-where-to-use-iam-permissions-boundaries/>

upvoted 1 times

  **alce2020** 1 year, 4 months ago

mi vote is for C as the right answer



upvoted 1 times

  **asfsdfsdf** 1 year, 5 months ago

Selected Answer: A

Only valid solution is A, for C or D you need to attach boundaries on all IAM roles/users not the account or the role itself.

upvoted 1 times

  **ele** 1 year, 5 months ago

Selected Answer: C

Between A and C, A looks good, but "SCPs affect only member accounts in the organization. They have no effect on users or roles in the management account."

C would do the work for all accounts in the Organization.

upvoted 1 times

  **qan1257** 1 year, 3 months ago

All the accounts are in the same organization in AWS Organizations.

upvoted 3 times

A company has an on-premises application that is written in Go. A DevOps engineer must move the application to AWS. The company's development team wants to enable blue/green deployments and perform A/B testing. Which solution will meet these requirements?

- A. Deploy the application on an Amazon EC2 instance, and create an AMI of the instance. Use the AMI to create an automatic scaling launch configuration that is used in an Auto Scaling group. Use Elastic Load Balancing to distribute traffic. When changes are made to the application, a new AMI will be created, which will initiate an EC2 instance refresh.
- B. Use Amazon Lightsail to deploy the application. Store the application in a zipped format in an Amazon S3 bucket. Use this zipped version to deploy new versions of the application to Lightsail. Use Lightsail deployment options to manage the deployment.
- C. Use AWS CodeArtifact to store the application code. Use AWS CodeDeploy to deploy the application to a fleet of Amazon EC2 instances. Use Elastic Load Balancing to distribute the traffic to the EC2 instances. When making changes to the application, upload a new version to CodeArtifact and create a new CodeDeploy deployment.
- D. Use AWS Elastic Beanstalk to host the application. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the application. Use Elastic Beanstalk to manage the deployment options.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **zijo** 6 months, 1 week ago

AWS Elastic Beanstalk deploy action can be used to deploy the application artifact from the S3 bucket to the green environment, which is the AWS cloud environment here.

upvoted 3 times

🗳️ 👤 **dzn** 6 months, 3 weeks ago

Selected Answer: D

Lightsail does not have built-in Blue/Green deployment capabilities like Elastic Beanstalk.

upvoted 2 times

🗳️ 👤 **z_inderjot** 8 months, 3 weeks ago

Selected Answer: D

D is undoubtedly is most correct one . But they should mention that , we are going to deploy application in different environment . Since deploying to the same environment just override the previous deployment . In order to meet the requirement of Blue / Green deployment we need two separate environment . Then we have two separate version running simultaneously and we can do DNS swapping to quickly shift traffic .

upvoted 4 times

🗳️ 👤 **Kiroo** 1 year, 1 month ago

Selected Answer: D

I was about to discard D because I was unsure if beanstalk supported GO (yes it does)

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.platforms.html>

So D is undoubtedly the best option to quickly move to cloud and to do blue green with an testing

upvoted 3 times

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

Selected Answer: D

I guess D is easiest option to orchestrate blue/green deployments and A/B testing in this case.

upvoted 2 times

🗳️ 👤 **rdoty** 1 year, 3 months ago

Selected Answer: D

D elastic beanstalk due to deployment options

upvoted 1 times

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

Maybe D, but it looks like an old approach. we need to use codebuild and codepipelines and Elastic beanstalk, but elastic beanstalk could be changed by AWS cloudformation.

upvoted 1 times

  **haazybanj** 1 year, 4 months ago

Selected Answer: D

D. Use AWS Elastic Beanstalk to host the application. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the application. Use Elastic Beanstalk to manage the deployment options.



AWS Elastic Beanstalk provides a platform for deploying web applications, which is well-suited for use cases that require blue/green deployments and A/B testing. Elastic Beanstalk can deploy applications written in a variety of programming languages and frameworks, including Go. Elastic Beanstalk supports blue/green deployments, which allow you to deploy a new version of your application to a separate environment before switching traffic to it. This enables you to perform A/B testing before fully rolling out a new version of your application. Elastic Beanstalk also allows you to manage the deployment options, including the deployment strategy, instance types, and autoscaling options.

upvoted 4 times

  **alce2020** 1 year, 4 months ago

D it is

upvoted 1 times

  **ele** 1 year, 5 months ago

Selected Answer: D

Elastic Beanstalk

upvoted 1 times

A developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing.


Occasionally, some application servers are being terminated after failing ELB HTTP health checks. The developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated.

How can log collection be automated?

- A. Use Auto Scaling lifecycle hooks to put instances in a Pending:Wait state. Create an Amazon CloudWatch alarm for EC2 Instance Terminate Successful and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- B. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an AWS Config rule for EC2 Instance-terminate Lifecycle Action and trigger a step function that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- C. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon CloudWatch subscription filter for EC2 Instance Terminate Successful and trigger a CloudWatch agent that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- D. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon EventBridge rule for EC2 Instance-terminate Lifecycle Action and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

Suggested Answer: D

Community vote distribution



madperro Highly Voted 1 year, 9 months ago

Selected Answer: D

D is the easiest solution.

upvoted 10 times

wikn Most Recent 3 months, 1 week ago

Selected Answer: C

why C is not correct

upvoted 1 times

hayjaykay 1 month ago

Think fleet of EC2s, think SSM (systems manager).

upvoted 1 times

Ravi_Bulusu 3 months, 3 weeks ago

Option A is the most efficient and straightforward approach to automate log collection and prevent premature termination of EC2 instances by using Auto Scaling lifecycle hooks, CloudWatch alarms, Lambda functions, and SSM to gather and store logs in Amazon S3 before the instance is terminated.

upvoted 1 times

Saudis 3 months, 4 weeks ago

it is D not C because CloudWatch agent can not invokes a script

upvoted 1 times

jamesf 7 months, 1 week ago

Selected Answer: D

D as the EC2 is Terminating and Cloudwatch Agent should be not running and cannot collect the logs

upvoted 1 times

Rahul369 8 months, 3 weeks ago

Selected Answer: C

It must be 'C' as CloudWatch Agent will push the logs to a particular CloudWatch log group.

upvoted 1 times

🗳️ 👤 **dzn** 1 year ago

Selected Answer: D

Terminating:Wait refers to a state in which an instance is determined to be terminated by the Auto Scaling group as part of the termination process and is temporarily put on hold before it is actually terminated. This state pauses the termination process and provides an opportunity to perform custom actions (logging, graceful shutdown, data backup, etc).

upvoted 2 times

🗳️ 👤 **hoazgagh** 11 months ago

why not C bro

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

D is correct: Using Eventbridge in combination with lambda is a common practice.

A: Cloudwatch alarm only alert, no action so it cannot trigger lambda (when this question came out, it could not)

B: AWS config rule cannot trigger a script.

C: cloudwatch agent itself does not have any direct action on the host but collecting logs

upvoted 4 times

🗳️ 👤 **Jaguaroooo** 1 year, 2 months ago

C is also a good choice in this question. Why? you need to have a CW agent installed on the hosts to be able to collect logs from the servers before termination.

upvoted 1 times

🗳️ 👤 **davdan99** 1 year, 2 months ago

I think we can't select C because it says that it invokes the cloudwatch agent after the EC2 instance is terminated. It can't collect the logs from terminated EC2 Instance.

upvoted 2 times

🗳️ 👤 **bcx** 1 year, 9 months ago

Selected Answer: D

D is the correct one IMHO.

ASG actions are not logged to cloudwatch logs to use a filter, and if so it would be complicated to extract the data. The canonical way is to rely in an EventBridge event.

upvoted 3 times

🗳️ 👤 **levster** 1 year, 9 months ago

D

"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://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-instance/>

upvoted 4 times

🗳️ 👤 **vherman** 1 year, 10 months ago

Selected Answer: D

D for sure 100%

upvoted 1 times

🗳️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: D

D. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon EventBridge rule for EC2 Instance-terminate Lifecycle Action and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

With this solution, you can use an Auto Scaling lifecycle hook to put instances in a wait state before termination. This provides an opportunity to collect logs before the instance is terminated. The solution can use an Amazon EventBridge rule for EC2 Instance-terminate Lifecycle Action to trigger an AWS Lambda function that will execute an SSM Run Command script. The script can collect logs and push them to Amazon S3 before completing the lifecycle action and allowing the instance to terminate. This solution provides a way to collect logs before instances are terminated, allowing for root cause analysis of issues.

upvoted 4 times

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

Selected Answer: D

D seems to be more relevant for this scenario
upvoted 1 times

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

Note that there is a similar question on Tutorial Dojo and the answer is to "trigger cloudwatch agent"
upvoted 1 times

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

read this link and you will understand that C is wrong option- <https://aws.amazon.com/blogs/infrastructure-and-automation/run-code-before-terminating-an-ec2-auto-scaling-instance/>
upvoted 1 times

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

Selected Answer: C

should be C
upvoted 2 times

🗨️ 👤 **beanxyz** 1 year, 6 months ago

No way. Cloudwatch subscription filter is normally used to send cloudwatch log to kinesis firehose stream so that it can be consumed by other tools such as Splunk. If you need to invoke a lambda, the easiest way is to use event rule.
upvoted 2 times

🗨️ 👤 **ele** 1 year, 11 months ago

Selected Answer: D

D sure
upvoted 1 times

A company has an organization in AWS Organizations. The organization includes workload accounts that contain enterprise applications. The company centrally manages users from an operations account. No users can be created in the workload accounts. The company recently added an operations team and must provide the operations team members with administrator access to each workload account. Which combination of actions will provide this access? (Choose three.)

- A. Create a SysAdmin role in the operations account. Attach the AdministratorAccess policy to the role. Modify the trust relationship to allow the sts:AssumeRole action from the workload accounts.
- B. Create a SysAdmin role in each workload account. Attach the AdministratorAccess policy to the role. Modify the trust relationship to allow the sts:AssumeRole action from the operations account.
- C. Create an Amazon Cognito identity pool in the operations account. Attach the SysAdmin role as an authenticated role.
- D. In the operations account, create an IAM user for each operations team member.
- E. In the operations account, create an IAM user group that is named SysAdmins. Add an IAM policy that allows the sts:AssumeRole action for the SysAdmin role in each workload account. Add all operations team members to the group.
- F. Create an Amazon Cognito user pool in the operations account. Create an Amazon Cognito user for each operations team member.

Suggested Answer: ABE

Community vote distribution

BDE (88%)

8%

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

Selected Answer: BDE

Any thing Cognito, safe to remove (it is only used for application identity management)

Step 1: Create each role in each workload account. Set trust relationship to only sts:AssumeRole via the operations user in operations account

Step 2: Self explanatory: whatever permission you needs once the user assumed the role

Step 3: Voila

upvoted 7 times

 **jamesf** Most Recent 7 months, 2 weeks ago


Selected Answer: BDE

BDE

Not A - Create SysAdmin role for workload accounts.

Not C F - No Cognito require.

upvoted 1 times

 **HarryLy** 9 months, 1 week ago

Selected Answer: BDE

Operation account:

- Need to create a role to assume role in workload account --> E

- Create a group of users can perform assume role --> D

workload account

- Need to create a role with have admin perssion for operation account assume -->B

upvoted 1 times

 **c3518fc** 11 months ago

Selected Answer: BEF

Not sure why everyone is saying BDE. Why would you create an IAM user for each member and also create for the group? Make it make sense

upvoted 2 times

 **4555894** 1 year ago

Selected Answer: BDE

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

upvoted 1 times

 **Vitalydt** 1 year ago

Selected Answer: BDE

EBD looks like the best choice

upvoted 1 times

🗨️ **dzn** 1 year ago

Selected Answer: BDE

sts:AssumeRole is one of the AWS Security Token Service (STS) actions used to obtain temporary security credentials and assume the role of another AWS account.

upvoted 1 times

🗨️ **thanhnv142** 1 year, 1 month ago

BDE: No cognito here.

-step 1: create role in workload accounts

-step 2: create IAM user for each member

-step 3: move all member to the group that has permission to assume the role in step 1

upvoted 3 times

🗨️ **madperro** 1 year, 9 months ago

Selected Answer: BDE

BDE seems to be right.

upvoted 2 times

🗨️ **rdoty** 1 year, 9 months ago

Selected Answer: BDE

def BDE cause role must be created in workload accounts and assumed by the operations account

upvoted 1 times

🗨️ **bcx** 1 year, 9 months ago

Selected Answer: BDE

Correct: BDE

Cognito has nothing to do with this, so C and F are wrong.

The roles must be created in the workload accounts and assumed from the operations account. So A is wrong.

upvoted 1 times

🗨️ **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: BDE

BDE seems the correct strategy

upvoted 3 times

🗨️ **Saga** 1 year, 10 months ago

Why do we need option A when question is asking access to workload account?

upvoted 1 times

🗨️ **alce2020** 1 year, 11 months ago

A,B,E it is

upvoted 1 times

🗨️ **ele** 1 year, 11 months ago

Selected Answer: BDE

BDE is right answer, nothing to do with cognito

upvoted 2 times

🗨️ **jqso234** 1 year, 11 months ago

Selected Answer: ABE

Options C, D, and F are incorrect because they do not provide a way for the operations team members to assume a role in the workload accounts, which is necessary to access the resources in those accounts.

upvoted 1 times

🗨️ **vvndx** 1 year, 10 months ago

Should be BDE, Why the need to create two roles?

upvoted 1 times

🗨️ **boledadian** 1 year, 11 months ago

BDE is correct

upvoted 1 times


A company has multiple accounts in an organization in AWS Organizations. The company's SecOps team needs to receive an Amazon Simple Notification Service (Amazon SNS) notification if any account in the organization turns off the Block Public Access feature on an Amazon S3 bucket. A DevOps engineer must implement this change without affecting the operation of any AWS accounts. The implementation must ensure that individual member accounts in the organization cannot turn off the notification.

Which solution will meet these requirements?

- A. Designate an account to be the delegated Amazon GuardDuty administrator account. Turn on GuardDuty for all accounts across the organization. In the GuardDuty administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. In the same account, create an Amazon EventBridge rule that uses an event pattern for GuardDuty findings and a target of the SNS topic.
- B. Create an AWS CloudFormation template that creates an SNS topic and subscribes the SecOps team's email address to the SNS topic. In the template, include an Amazon EventBridge rule that uses an event pattern of CloudTrail activity for s3:PutBucketPublicAccessBlock and a target of the SNS topic. Deploy the stack to every account in the organization by using CloudFormation StackSets.
- C. Turn on AWS Config across the organization. In the delegated administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. Deploy a conformance pack that uses the s3-bucket-level-public-access-prohibited AWS Config managed rule in each account and uses an AWS Systems Manager document to publish an event to the SNS topic to notify the SecOps team.
- D. Turn on Amazon Inspector across the organization. In the Amazon Inspector delegated administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. In the same account, create an Amazon EventBridge rule that uses an event pattern for public network exposure of the S3 bucket and publishes an event to the SNS topic to notify the SecOps team.

Suggested Answer: B

Community vote distribution



rif Highly Voted 1 year, 4 months ago

Answer is C.

* AWS Systems Manager Automation provides predefined runbooks(ex. AWS-PublishSNSNotification) for Amazon Simple Notification Service - <https://docs.aws.amazon.com/systems-manager-automation-runbooks/latest/userguide/automation-aws-publishsnsnotification.html>

* Running automations in multiple AWS Regions and accounts (<https://docs.aws.amazon.com/systems-manager/latest/userguide/running-automations-multiple-accounts-regions.html>)

B seems to be old approach. With cloudformation stackset, each account can still change resource config (ex. SNS) that causes drift.... so I choose C because it utilize AWS organization fully with aws systems manager automation in multiple regions and multiple accounts with delegated administrator account(or management account)

upvoted 11 times

flacko 6 months, 3 weeks ago

With option B, you will get notifications when user accounts turn off the block public access feature but it doesn't stop them from doing it. The question requires that the implementation stops users from being able to carry out that operation successfully altogether.

upvoted 2 times

Impromptu 3 months, 3 weeks ago

Just to go more into detail, as the answer C seems correct indeed. But I'd like to point out some extra details on why B is wrong.

The question asks that a user cannot turn off the notification. They should be able to turn off the block public access feature. So B is not wrong because it doesn't implement the latter.

B is wrong because it's PutPublicAccessBlock (does not contain "Bucket" in it). And additionally, you should add a condition to the eventbridge rule that checks the content of the action: that BlockPublicPolicy is set to False. Without the condition you will get notification on all PutPublicAccessBlock events, so also those that are considered to be valid.

upvoted 1 times

Impromptu 3 months, 3 weeks ago

To bad I can't edit, so to correct myself: PutBucketPublicAccessBlock is indeed the IAM permission and what you should filter on.

And the cloudformation solution in option B also lacks the safeguard to prevent users from disabling the eventbridge rule (and therefore disabling the notification)

upvoted 1 times

🗨️ **Gomer** Most Recent 6 months ago

Selected Answer: A

GuardDuty Policy

Policy:S3/BucketBlockPublicAccessDisabled

"An IAM entity invoked an API used to disable S3 Block Public Access on a bucket."

"Data source: CloudTrail management events"

"This finding informs you that Block Public Access was disabled for the listed S3 bucket. When enabled, S3 Block Public Access settings are used to filter the policies or access control lists (ACLs) applied to buckets as a security measure to prevent inadvertent public exposure of data."

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_finding-types-s3.html#policy-s3-bucketblockpublicaccessdisabled

upvoted 1 times

🗨️ **jamesf** 7 months, 2 weeks ago

Selected Answer: C

C

"A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations."

<https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html>

<https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>

upvoted 1 times

🗨️ **aefuen1** 8 months, 1 week ago

Selected Answer: A

It's A. GuardDuty achieves this with no effort.

upvoted 1 times

🗨️ **xdkonorek2** 8 months, 2 weeks ago

Selected Answer: A

A DevOps engineer must implement this change without affecting the operation of any AWS accounts.

upvoted 2 times

🗨️ **Gomer** 9 months, 2 weeks ago

I was sure the answer was "C" until I started reading through some of the requirements and comments. The words "implementation must ensure that individual member accounts in the organization cannot turn off the notification" incline me to lean towards "A", because with "C", someone with admin privileges on a single account could turn off the notification in that account. As pointed out by others, there are a number of GuardDuty findings associates with S3 public access. Having GuardDuty and EventBridge pattern trigger SNS for some key words such as "s3" and "Public" seems to make sense in enforcing this across an organization. I don't have enough experience with GuardDuty in an Organization to be 100% confident, but the emphasis on SNS requirement makes me think this could be a trick question.

upvoted 1 times

🗨️ **Gomer** 6 months ago

GuardDuty Policy

Policy:S3/BucketBlockPublicAccessDisabled

"An IAM entity invoked an API used to disable S3 Block Public Access on a bucket."

"Data source: CloudTrail management events"

"This finding informs you that Block Public Access was disabled for the listed S3 bucket. When enabled, S3 Block Public Access settings are used to filter the policies or access control lists (ACLs) applied to buckets as a security measure to prevent inadvertent public exposure of data."

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_finding-types-s3.html#policy-s3-bucketblockpublicaccessdisabled

upvoted 1 times

🗨️ **seetpt** 10 months, 2 weeks ago

Selected Answer: C

C is only correct option.

upvoted 1 times

🗨️ **that1guy** 10 months, 3 weeks ago

Technically A would be sufficient here.

The question is only asking to be NOTIFIED when block public access gets disabled.

See the following GuardDuty finding: https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_finding-types-s3.html#policy-s3-

bucketblockpublicaccessdisabled

Managing multiple GuardDuty accounts is simplified using the AWS Organizations delegated administrator feature. With this feature, the AWS Organizations management account can designate a member account to be the GuardDuty administrator for the entire organization. The delegated GuardDuty administrator is then granted permission to enable and manage GuardDuty for all existing and future accounts in the organization.

upvoted 3 times

🗳️ 👤 **Cervus18** 12 months ago

Selected Answer: A

We can leverage AWS Organizations to enable Guardduty in all accounts.

There is an S3 finding called Policy:S3/AccountBlockPublicAccessDisabled

Then we setup a single EventBrdige rule in the delegated account that publish the event to the SNS topic in the same account.

This is the easisest solution to be implemented and monitoring the public access seamlessly across all Organization's accounts

This is a common multi-account strategy for GuardDuty with AWS organizations, to collect such finding from hundred of accounts

upvoted 4 times

🗳️ 👤 **4555894** 1 year ago

Selected Answer: C

Amazon GuardDuty is primarily on threat detection and response, not configuration monitoring. A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations. <https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html><https://docs.aws>.

upvoted 4 times

🗳️ 👤 **zijo** 1 year ago

Answer is C

A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations. You can also use AWS Systems Manager documents (SSM documents) to store your conformance pack templates on AWS and directly deploy conformance packs using SSM document names.

upvoted 3 times

🗳️ 👤 **Rocky007** 1 year ago

Hi can somebody with contributors access, would please forward all the questions pdf to me on telegram @rater250 , I'm willing to pay

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

C is correct: AWS config can only be modify by admin, not member accounts

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

Let me clarify: B cannot be correct because of this reason: "Deploy the stack to every account in the organization by using CloudFormation StackSets" means in every accounts of this AWS org (canbe up to hundreds of account), we will deploy a SNS topic and an EventBridge rule. This would be an extremely expensive deployment

upvoted 2 times

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

Option B is also not a valid case because we can direct use config with eventbrige why to go for clod trail we can use aws config rule s3-bucket-public-read-prohibited if rule changes eventbridge will trigger sns

upvoted 2 times

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

I got confused with option B and C , but Lets think in C option when I will use system manager to trigger SNS I can simply use eventbridge run that checks for config rule compliance change , IF compliance changes then as a target we will specify SNS.

Yes , We can also specify system manager automation document to trigger sns but why I will use it I will directly use SNS.

So from above I still by looking words B is correct option. Main reason is you do not need system manager here to trigger SNS.

Plus there is no mention for eventbridge rule that will trigger system manager , from config we cannot directly trigger it.

upvoted 1 times

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

I got confused with option B and C , but Lets think in C option when I will use system manager to trigger SNS I can simply use eventbridge run that checks for config rule compliance change , IF compliance changes then as a target we will specify SNS.

Yes , We can also specify system manager automation document to trigger sns but why I will use it I will directly use SNS.

So from above I still by looking words B is correct option. Main reason is you do not need system manager here to trigger SNS.

upvoted 1 times

  **Jay_2pt0_1** 1 year, 2 months ago

Selected Answer: C

This is the type of thing that AWS Config is used for.

upvoted 2 times

A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EKS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic. Which logging solution will support these requirements?

- A. Enable Amazon CloudWatch Logs to log the EKS components. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- B. Enable Amazon CloudWatch Logs to log the EKS components. Create CloudWatch Logs Insights queries linked to Amazon EventBridge events that invoke Lambda.
- C. Enable Amazon S3 logging for the EKS components. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- D. Enable Amazon S3 logging for the EKS components. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

Suggested Answer: C

Community vote distribution

A (97%)

 **tartarus23** Highly Voted 1 year, 2 months ago

Selected Answer: A

Correct Answer is A.

Explanation:

Amazon EKS integrates with CloudWatch Logs to provide detailed logs of the state and execution of the services in the cluster. CloudWatch subscription filters can be used to route specific log events from a CloudWatch Logs group to a Lambda function. The Lambda function can then process the events and publish notifications to the appropriate Amazon SNS topic.

upvoted 15 times

 **4555894** Most Recent 6 months ago

Selected Answer: A

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html#LambdaFunctionExample>

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>

upvoted 3 times

 **zijo** 6 months, 1 week ago

AWS EKS itself does not offer native S3 logging for container logs. CloudWatch Logs Insights queries cannot directly link to Amazon EventBridge events. So the answer here is A.

upvoted 1 times

 **dzn** 6 months, 3 weeks ago

Selected Answer: A

CloudWatch Logs subscription filtering is a feature that allows capture log data in real time and forward it to other AWS services such as Kinesis Data Firehose, Kinesis Streams, and Lambda.

upvoted 1 times


 **thanhnv142** 7 months, 2 weeks ago

A is correct: Use cloudwatch logs to collect logs from EKS. Use subscription filter to filter out logs and only send relevant logs to lambda to trigger it.

B: CloudWatch Logs Insights is for data analysis. Additionally, using EventBridge events to trigger lambda incur costs

C and D: Amazon S3 logging is used for monitoring actions on S3 itself, not EKS

upvoted 3 times

 **z_inderjot** 8 months, 3 weeks ago

Selected Answer: A

A is right

C, D are wrong , because there is not integration in EKS to send logs to s3.

B is for log analysis , and aggregation

upvoted 4 times

🗨️ 👤 **zolphar_z** 9 months, 3 weeks ago

Selected Answer: B

I don't have a technical reason but others dumps shows B as the Answer
upvoted 1 times

🗨️ 👤 **zolphar_z** 9 months, 3 weeks ago

No, sorry, this was for the previous questions
upvoted 2 times

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

Selected Answer: A

A, metric filter can call Lambda.
upvoted 1 times

🗨️ 👤 **rdoty** 1 year, 3 months ago

Selected Answer: A

certainly cloudwatch logs metric filter A
upvoted 1 times

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

Selected Answer: A

A. Enable Amazon CloudWatch Logs to log the EKS components. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.

This solution involves enabling Amazon CloudWatch Logs to log the EKS components and creating a CloudWatch subscription filter for each component with AWS Lambda as the subscription feed destination. This approach will allow the Lambda function to evaluate incoming log events and publish messages to the correct Amazon SNS topic. Amazon SNS can then send email notifications to each email address based on the messages it receives from the corresponding SNS topic.

upvoted 1 times

🗨️ 👤 **ele** 1 year, 4 months ago

Selected Answer: A

A, clear
upvoted 1 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

A is the correct answer
upvoted 1 times

🗨️ 👤 **jqso234** 1 year, 4 months ago

Selected Answer: A

Amazon CloudWatch Logs can log the EKS components, and subscription filters can be created for each component with AWS Lambda as the subscription feed destination. The Lambda function can evaluate incoming log events and publish messages to the appropriate Amazon SNS topic, enabling automated email notifications to be sent. Therefore, option A is the correct solution. Option C is incorrect because Amazon S3 logging is not designed for logging EKS components.

upvoted 1 times

🗨️ 👤 **Dimidrol** 1 year, 5 months ago

Selected Answer: A

A for sure
upvoted 3 times

A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic.

A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.


Which combination of steps must the DevOps engineer take to meet these requirements? (Choose three.)

- A. Download the Amazon CloudWatch Logs container instance from AWS. Configure this instance as a task. Update the application service definitions to include the logging task.
- B. Install the Amazon CloudWatch Logs agent on the ECS instances. Change the logging driver in the ECS task definition to awslogs.
- C. Use Amazon EventBridge to schedule an AWS Lambda function that will run every 60 seconds and will run the Amazon CloudWatch Logs create-export-task command. Then point the output to the logging S3 bucket.
- D. Activate access logging on the ALB. Then point the ALB directly to the logging S3 bucket.
- E. Activate access logging on the target groups that the ECS services use. Then send the logs directly to the logging S3 bucket.
- F. Create an Amazon Kinesis Data Firehose delivery stream that has a destination of the logging S3 bucket. Then create an Amazon CloudWatch Logs subscription filter for Kinesis Data Firehose.

Suggested Answer: BDE

Community vote distribution

BDF (100%)

 **tartarus23** Highly Voted 1 year, 8 months ago

Selected Answer: BDF

Explanation:

Option B is correct because you can change the logging driver in the ECS task definition to awslogs, which will direct the logs to Amazon CloudWatch Logs. Then, the logs can be forwarded to the Amazon S3 bucket.

Option D is correct because enabling access logging on the Application Load Balancer (ALB) allows the collection of access logs that can be sent directly to an S3 bucket.

Option F is correct because you can create an Amazon Kinesis Data Firehose delivery stream that can deliver logs from CloudWatch Logs directly to an Amazon S3 bucket in near-real-time.

upvoted 14 times

 **steli0** Most Recent 3 months, 2 weeks ago

Selected Answer: BDF

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html>

upvoted 2 times

 **4555894** 1 year ago

Selected Answer: BDF

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html>

upvoted 2 times

 **dzn** 1 year ago

Selected Answer: BDF

Enable access logging using the ALB management console, CLI, or API. Specify the S3 bucket where the logs will be stored and, if necessary, set the log file prefix (e.g., production, staging.) to store the logs in different paths within the bucket.

upvoted 1 times

 **thanhnv142** 1 year, 1 month ago

BDF: There are two types of log that needs to be collected

B: push app log to Cloudwatch log

D: push access log to S3

F: using Kinesis to push app log from cloudwatch log to S3 in near real-time

A: wrong - we need cloudwatch agent, not container instance

C: No need to use event bridge and lambda to trigger cloudwatch log to push log to s3.

E: access logs lie in ALB, not ECS services.

upvoted 4 times

🗨️ **z_inderjot** 1 year, 2 months ago

Selected Answer: BDF

BDF is the answer .

btw, can't we use cloudwatch ingests to collect the logs from containers in ecs there days , and then use the subscription filter we can send those logs to s3.

without having to install cloud watch agent.

upvoted 2 times

🗨️ **imymoco** 1 year, 3 months ago

Real time. so not E

upvoted 1 times

🗨️ **madperro** 1 year, 8 months ago

Selected Answer: BDF

BDF makes sense. E is certainly wrong.

upvoted 1 times

🗨️ **bcx** 1 year, 9 months ago

Selected Answer: BDF

BDF

Access logs cannot be configured by ALB target group

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

upvoted 2 times

🗨️ **hanbj** 1 year, 9 months ago

Option B sends data to the Cloudwatch Log. This issue requires that logs be collected in S3.

upvoted 1 times

🗨️ **haazybanj** 1 year, 10 months ago

Selected Answer: BDF

Answer is bdf

upvoted 2 times

🗨️ **ele** 1 year, 11 months ago

Selected Answer: BDF

B - get application logs to CW

D - get access logs to S3

F - get application logs from CW to S3 in near-real time

upvoted 4 times

🗨️ **jqso234** 1 year, 11 months ago

Selected Answer: BDF

Option BDE can be cumbersome to manage in a large environment and may not be ideal for applications that generate large amounts of logs.

Option BDF, on the other hand, captures both application and access logs, and uses the CloudWatch Logs driver to stream logs directly to CloudWatch Logs. This solution is more scalable as it does not require the CloudWatch Logs agent to be installed on each instance, and it can capture logs from multiple ECS tasks running on the same instance. In addition, the logs can be sent to an S3 bucket using a Kinesis Data Firehose delivery stream, which provides near-real-time analysis capabilities.

upvoted 1 times

🗨️ **Dimidrol** 1 year, 11 months ago

Selected Answer: BDF

B D F for me

upvoted 2 times

🗨️ **Dimidrol** 1 year, 11 months ago

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/using_cloudwatch_logs.html

upvoted 1 times

A company that uses electronic health records is running a fleet of Amazon EC2 instances with an Amazon Linux operating system. As part of patient privacy requirements, the company must ensure continuous compliance for patches for operating system and applications running on the EC2 instances.

How can the deployments of the operating system and application patches be automated using a default and custom repository?

- A. Use AWS Systems Manager to create a new patch baseline including the custom repository. Run the AWS-RunPatchBaseline document using the run command to verify and install patches.
- B. Use AWS Direct Connect to integrate the corporate repository and deploy the patches using Amazon CloudWatch scheduled events, then use the CloudWatch dashboard to create reports.
- C. Use yum-config-manager to add the custom repository under /etc/yum.repos.d and run yum-config-manager-enable to activate the repository.
- D. Use AWS Systems Manager to create a new patch baseline including the corporate repository. Run the AWS-AmazonLinuxDefaultPatchBaseline document using the run command to verify and install patches.

Suggested Answer: A

Community vote distribution

A (100%)


  **dzn**  6 months, 3 weeks ago

Selected Answer: A

AWS-AmazonLinuxDefaultPatchBaseline: defines which patches should be applied and which should be avoided.

AWS-RunPatchBaseline: provides commands to actually run the patching process on the instance.

upvoted 6 times



  **thanhnv142**  7 months, 2 weeks ago

A is correct: AWS system manager and AWS-RunPatchBaseline to utilize a default and custom repo

B and C are irrelevant

D: AWS-AmazonLinuxDefaultPatchBaseline: this baseline has "default" in its name, it is a predefined baseline and cannot work with a custom repo

upvoted 4 times

  **davdan99** 8 months, 1 week ago



Here are predefined documents that can not be modified (includes AWS-AmazonLinuxDefaultPatchBaseline)

<https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-predefined-and-custom-patch-baselines.html#patch-manager-baselines-custom>

And here is about the AWS-RunPatchBaseline



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

upvoted 2 times

  **davdan99** 8 months, 1 week ago

the Answer is A

upvoted 1 times

  **z_inderjot** 8 months, 3 weeks ago

Selected Answer: A

I was confused between A and D , i choose A instinctively, D statement sounds like it only going to install default package for linux not from custom repo we add . But not sure any one can clarify


upvoted 3 times

  **madperro** 1 year, 2 months ago

Selected Answer: A

A, SSM allows inclusion of custom repositories.

upvoted 4 times

  **haazybanj** 1 year, 4 months ago

Selected Answer: A

A is it



upvoted 1 times

  **alce2020** 1 year, 4 months ago

Selected Answer: A

A is correct

upvoted 2 times

  **jqso234** 1 year, 4 months ago

Selected Answer: A

To automate the deployment of operating system and application patches using a default and custom repository in Amazon EC2 instances with Amazon Linux operating systems, you can use AWS Systems Manager. You can create a new patch baseline in Systems Manager that includes the custom repository, then run the AWS-RunPatchBaseline document using the run command to verify and install patches. This allows you to ensure continuous compliance for patches while also automating the patch deployment process.

upvoted 1 times

A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon Elastic Container Service (Amazon ECS) using the blue/green deployment model. The company wants to implement scripts to test the green version of the application before shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back.



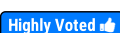
Which strategy will meet these requirements?

- A. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use AWS CodeBuild to create a runtime environment and build commands in the buildspec file to invoke test scripts. If errors are found, use the aws deploy stop-deployment command to stop the deployment.
- B. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use this stage to invoke an AWS Lambda function that will run the test scripts. If errors are found, use the aws deploy stop-deployment command to stop the deployment.
- C. Add a hooks section to the CodeDeploy AppSpec file. Use the AfterAllowTestTraffic lifecycle event to invoke an AWS Lambda function to run the test scripts. If errors are found, exit the Lambda function with an error to initiate rollback.
- D. Add a hooks section to the CodeDeploy AppSpec file. Use the AfterAllowTraffic lifecycle event to invoke the test scripts. If errors are found, use the aws deploy stop-deployment CLI command to stop the deployment.

Suggested Answer: C

Community vote distribution

C (100%)




  **haazybanj**  1 year, 4 months ago

Add a hooks section to the CodeDeploy AppSpec file. The AppSpec file is a YAML file that describes how to deploy an application to Amazon ECS using CodeDeploy. We can use the AfterAllowTestTraffic lifecycle event to run the test scripts. This event is triggered after the new version of the application is deployed, and before traffic is shifted to the new version.

In the AfterAllowTestTraffic lifecycle event, invoke an AWS Lambda function to run the test scripts. The Lambda function can be written in any programming language supported by Lambda, such as Python, Node.js, or Java.

If the test scripts detect any errors, exit the Lambda function with an error code. This will cause the deployment to fail, and CodeDeploy will initiate a rollback.

upvoted 10 times

  **zijo**  6 months ago

AfterAllowTestTraffic lifecycle event in the hooks section will not shift the whole traffic to the green application but only a small percentage of traffic to the newly deployed version. C is the answer

upvoted 1 times

  **dzn** 6 months, 3 weeks ago

Selected Answer: C

CodeDeploy Blue/Green deployments, the AfterAllowTestTraffic hook is triggered after the test traffic redirection to the new version (Green) is set. Additional verification, testing, or other custom actions can be automated by executing Lambda functions at this time.

upvoted 2 times

  **thanhnv142** 7 months, 2 weeks ago

C is correct: we can initiate the script using lambda for advanced features

A and B are wrong: Both trigger the test script before deploy stages

D is wrong: It only stops the deployment, not rollback it

upvoted 4 times

  **ixdb** 1 year ago

C is right.

upvoted 1 times

  **madperro** 1 year, 2 months ago

Selected Answer: C

C is the right answer.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-ecs>

upvoted 4 times

  **haazybanj** 1 year, 4 months ago

Selected Answer: C

The correct solution to meet these requirements is option C.

Explanation:

In this scenario, the requirement is to add scripts to test the green version of the application before shifting traffic. These scripts should be executed quickly and, in case of errors, the application must be rolled back. To achieve this, we can use the following steps:

upvoted 1 times


  **ele** 1 year, 4 months ago

Selected Answer: C

Lifecycle event hooks for an Amazon ECS deployment:

AfterAllowTraffic – Use to run tasks after the second target group serves traffic to the replacement task set. The results of a hook function at this lifecycle event can trigger a rollback.

upvoted 1 times

  **ele** 1 year, 4 months ago

Correction:

AfterAllowTestTraffic – Use to run tasks after the test listener serves traffic to the replacement task set. The results of a hook function at this point can trigger a rollback.

upvoted 1 times

  **alice2020** 1 year, 4 months ago

Selected Answer: C

C is the correct answer

upvoted 1 times

A company uses AWS Storage Gateway in file gateway mode in front of an Amazon S3 bucket that is used by multiple resources. In the morning when business begins, users do not see the objects processed by a third party the previous evening. When a DevOps engineer looks directly at the S3 bucket, the data is there, but it is missing in Storage Gateway.



Which solution ensures that all the updated third-party files are available in the morning?

- A. Configure a nightly Amazon EventBridge event to invoke an AWS Lambda function to run the RefreshCache command for Storage Gateway.
- B. Instruct the third party to put data into the S3 bucket using AWS Transfer for SFTP.
- C. Modify Storage Gateway to run in volume gateway mode.
- D. Use S3 Same-Region Replication to replicate any changes made directly in the S3 bucket to Storage Gateway.

Suggested Answer: A

Community vote distribution

A (97%)

  **tartarus23** Highly Voted 1 year, 2 months ago

Selected Answer: A

Explanation:

AWS Storage Gateway's file gateway mode provides a bridge between your on-premises servers and Amazon S3. File gateway caches frequently accessed files in your on-premises environment to provide low-latency access. However, if the S3 bucket's data is modified by another service, the cache does not automatically refresh. Thus, to ensure all the updated third-party files are available in the morning, you can use an AWS Lambda function triggered by Amazon EventBridge to run the RefreshCache command for Storage Gateway. This will ensure the cache is updated with the latest changes.

upvoted 16 times

  **ele** Highly Voted 1 year, 4 months ago

Selected Answer: A

A: refresh cache: <https://repost.aws/knowledge-center/storage-gateway-s3-changes-not-showing>

upvoted 12 times

  **robertohyena** 10 months ago

Thanks for this.

Also found <https://repost.aws/knowledge-center/storage-gateway-automate-refreshcache>

Storage Gateway allows you to automate the RefreshCache operation based on a Time To Live (TTL) value. TTL is the length of time since the last refresh. When a user accesses the file directory after the TTL value, the file gateway refreshes the directory's contents from the S3 bucket. Valid TTL values for automating the RefreshCache operation range from 300 seconds to 2,592,000 seconds (5 minutes to 30 days).

upvoted 3 times

  **Gomer** Most Recent 3 months, 1 week ago

Selected Answer: A

Read and concede:

"Configure an automated cache refresh schedule using AWS Lambda with an Amazon CloudWatch rule"



<https://docs.aws.amazon.com/filegateway/latest/files3/refresh-cache.html#auto-refresh-lambda-procedure>

upvoted 1 times

  **bhond** 1 year ago

where is it saying files are written directly to s3 ?

upvoted 1 times

  **yorkicurve** 9 months, 3 weeks ago

You do make a point but if you read the phrase " When a DevOps engineer looks directly at the S3 bucket " it kinda implies besides you dont have any other better choice anyway. if you look at user "ele" comments and follow the link below it will get clear[hope that helps];

<https://repost.aws/knowledge-center/storage-gateway-s3-changes-not-showing>

upvoted 3 times

🗨️ 👤 **ixdb** 1 year ago

A is right.

Storage Gateway updates the file share cache automatically when you write files to the cache locally using the file share. However, Storage Gateway doesn't automatically update the cache when you upload a file directly to Amazon S3. When you do this, you must perform a RefreshCache operation to see the changes on the file share.

upvoted 1 times

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

Selected Answer: A

A is the answer.

upvoted 1 times

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

Selected Answer: A

The issue appears to be related to the Storage Gateway cache not being updated. To ensure that all the updated third-party files are available in the morning, you can use the RefreshCache API to manually refresh the cache or configure automatic cache refresh.

Option A is a possible solution to configure automatic cache refresh, but it is not necessary to run the RefreshCache command every night if you can ensure that cache refresh occurs frequently enough to meet your requirements.

upvoted 2 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

Selected Answer: A

A is correct

upvoted 1 times

🗨️ 👤 **jqso234** 1 year, 4 months ago

Selected Answer: B

Option B appears to be the correct choice. Configuring the third party to put data into the S3 bucket using AWS Transfer for SFTP would ensure that the data is immediately available in both the S3 bucket and Storage Gateway, avoiding any potential caching issues. Option A of configuring a nightly event to refresh the cache may not be an optimal solution as it could result in stale data being served during the day.

upvoted 1 times

🗨️ 👤 **bcx** 1 year, 3 months ago

Transfer SFTP has the same effect in this case as adding files to S3 with PutObject. The cache in the storage gateway would not be updated, requiring the same refresh as in option A.

upvoted 1 times



A DevOps engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using S3 cross-Region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account. Which combination of actions should be performed to enable this replication? (Choose three.)

- A. Create a replication IAM role in the source account
- B. Create a replication IAM role in the target account.
- C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects.
- D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.
- E. Create a replication rule in the source bucket to enable the replication.
- F. Create a replication rule in the target bucket to enable the replication.

Suggested Answer: ADE

Community vote distribution

ADE (100%)

 **tschenhau**  1 year, 4 months ago

Selected Answer: ADE

S3 cross-Region replication (CRR) automatically replicates data between buckets across different AWS Regions. To enable CRR, you need to add a replication configuration to your source bucket that specifies the destination bucket, the IAM role, and the encryption type (optional). You also need to grant permissions to the IAM role to perform replication actions on both the source and destination buckets. Additionally, you can choose the destination storage class and enable additional replication options such as S3 Replication Time Control (S3 RTC) or S3 Batch Replication.

upvoted 7 times

 **Gomer**  3 months, 1 week ago

Selected Answer: ADE

Tricky question because they are trying to get one to confuse the "enable" replicaton "role"/policy ("rule") in source account with the "allow" replicaton role/"policy" in target account. These references helped me work up some summary steps:

Steps to configure S3 replication between different accounts

1. Create source and destination buckets in different accounts and regions (acctA, acctB)
2. Enable versioning on the buckets (acctA, acctB)
3. Create IAM role and attach a policy granting S3 permission to replicate objects (acctA)
4. Add the replication configuration to source bucket (acctA)
5. Add bucket "policy on the destination bucket to allow" objects replication (acctB)(req. 2nd role)

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough1.html#enable-replication>

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-2.html>

upvoted 3 times

 **Gomer** 3 months, 1 week ago

Source Account: (Source Bucket)(Versioning)(Role/Policy to "Enable" Replicaton)

Target Account: (Target Bucket)(Versioning)(Role/Policy to "Allow" Replicaton)

upvoted 2 times

 **thanhnv142** 7 months, 2 weeks ago

ADF is correct: this task is done by S3 itself

A: Create role in the source to allow S3 access permission

D: add policy to allow replication in the target

E: enable replication in the source

upvoted 2 times

 **bugincloud** 11 months, 3 weeks ago

Selected Answer: ADE

ADE make sense.

upvoted 2 times

 **madperro** 1 year, 2 months ago

Selected Answer: ADE

ADE make sense.


upvoted 2 times

  **haazybanj** 1 year, 4 months ago

Selected Answer: ADE



Confirmed

upvoted 2 times

  **alice2020** 1 year, 4 months ago

ADE confirmed!

upvoted 1 times

  **ele** 1 year, 4 months ago

Selected Answer: ADE

ADE

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-walkthrough-2.html>

upvoted 3 times


A company has multiple member accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the member accounts using an AWS Lambda function in the management account of the organization.

Which combination of access changes will meet these requirements? (Choose three.)

- A. Create a trust relationship that allows users in the member accounts to assume the management account IAM role.
- B. Create a trust relationship that allows users in the management account to assume the IAM roles of the member accounts.
- C. Create an IAM role in each member account that has access to the AmazonEC2ReadOnlyAccess managed policy.
- D. Create an IAM role in each member account to allow the sts:AssumeRole action against the management account IAM role's ARN.
- E. Create an IAM role in the management account that allows the sts:AssumeRole action against the member account IAM role's ARN.
- F. Create an IAM role in the management account that has access to the AmazonEC2ReadOnlyAccess managed policy.

Suggested Answer: BCE

Community vote distribution



tartarus23 Highly Voted 1 year, 2 months ago

Selected Answer: BCE

Explanation:

(B) The trust relationship enables an IAM entity (user, group, or role) to assume a role. In this case, the entities in the management account need to assume roles in the member accounts.

(C) The IAM role in each member account should have a policy attached that grants read-only access to EC2 instances. The AmazonEC2ReadOnlyAccess managed policy provides this access.

(E) An IAM role in the management account should be created that has the permission to perform the sts:AssumeRole action against the member account IAM role's ARN. This allows entities assuming this role to switch to the roles in the member accounts and perform actions according to the permissions of those roles.

upvoted 8 times

thanhnv142 Most Recent 7 months, 2 weeks ago

BCE are correct:

B: create trust relationship for management to assume role in member accounts

C: create role in member account that has access to AmazonEC2

E: Create IAM role in management account that allow access to member account IAM role

upvoted 4 times

svjl 9 months, 2 weeks ago

The security team wants to programmatically retrieve this information from the member accounts using an AWS Lambda function in the management account of the organization.

ReadOnlyAccess and option B grant the assumeRole

Besides that the correct resource is "IAM" not "I AM" So BCE is correct

upvoted 1 times

RVivek 1 year ago

Selected Answer: BCE

B- Member accounts should trust Management account



C- Member accounts should have a Role that has the necessary permission

E- Management account should have a IAM user account that has stsAssume role permission for the roles created in member accounts

upvoted 1 times

incorrigible_maverick 1 year ago

BCE is wrong. They want to programmatically therefore B is definitely wrong. The Lambda function IAM Role ARN in the management account needs to be able to assume a role in the member account that has the AmazonEC2ReadOnlyAccess attached to it. Therefore, I will go with C, D, E
upvoted 2 times

  **zain1258** 10 months, 1 week ago

D is clearly wrong. You are running your lambda function to get details in management account. The IAM role should be in management account with sts:AssumeRole permission to assume IAM roles in member accounts
upvoted 1 times

  **DavidPham** 1 year, 1 month ago



BCE correct
upvoted 1 times

  **madperro** 1 year, 2 months ago


Selected Answer: BCE
BCE is right.
upvoted 1 times

  **bcx** 1 year, 3 months ago

B, C and E
upvoted 1 times

  **PhuocT** 1 year, 3 months ago



Selected Answer: BCE
B, C and E
upvoted 2 times

  **2pk** 1 year, 3 months ago

Selected Answer: ACE
A:By creating a trust relationship that allows users in the member accounts to assume the IAM role in the management account, they will have the necessary permissions to access resources and retrieve the required information.

C:To grant the necessary permissions for retrieving information about EC2 security groups, an IAM role should be created in each member account. This role should have the AmazonEC2ReadOnlyAccess managed policy attached, which provides the required permissions.

E:In the management account, an IAM role should be created that allows assuming the IAM role in the member accounts. This role should have the necessary permissions to perform the sts:AssumeRole action against the ARN of the IAM roles in the member accounts.
upvoted 2 times

  **ele** 1 year, 3 months ago


Selected Answer: BCE
BCE will create correct cross account permission
upvoted 1 times

  **vherman** 1 year, 4 months ago

Selected Answer: ACD
acd looks good)
upvoted 1 times

  **marcoforexam** 1 year, 4 months ago

Selected Answer: ACD
ACE I guess
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html
upvoted 1 times

  **haazybanj** 1 year, 4 months ago

Selected Answer: BCE
Correct answer
upvoted 1 times

  **alce2020** 1 year, 4 months ago

Selected Answer: BCE
Ill go with BCF
upvoted 2 times

A space exploration company receives telemetry data from multiple satellites. Small packets of data are received through Amazon API Gateway and are placed directly into an Amazon Simple Queue Service (Amazon SQS) standard queue. A custom application is subscribed to the queue and transforms the data into a standard format.

Because of inconsistencies in the data that the satellites produce, the application is occasionally unable to transform the data. In these cases, the messages remain in the SQS queue. A DevOps engineer must develop a solution that retains the failed messages and makes them available to scientists for review and future processing.

Which solution will meet these requirements?

- A. Configure AWS Lambda to poll the SQS queue and invoke a Lambda function to check whether the queue messages are valid. If validation fails, send a copy of the data that is not valid to an Amazon S3 bucket so that the scientists can review and correct the data. When the data is corrected, amend the message in the SQS queue by using a replay Lambda function with the corrected data.
- B. Convert the SQS standard queue to an SQS FIFO queue. Configure AWS Lambda to poll the SQS queue every 10 minutes by using an Amazon EventBridge schedule. Invoke the Lambda function to identify any messages with a `SentTimestamp` value that is older than 5 minutes, push the data to the same location as the application's output location, and remove the messages from the queue.
- C. Create an SQS dead-letter queue. Modify the existing queue by including a redrive policy that sets the `MaximumReceives` setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue. Instruct the scientists to use the dead-letter queue to review the data that is not valid. Reprocess this data at a later time.
- D. Configure API Gateway to send messages to different SQS virtual queues that are named for each of the satellites. Update the application to use a new virtual queue for any data that it cannot transform, and send the message to the new virtual queue. Instruct the scientists to use the virtual queue to review the data that is not valid. Reprocess this data at a later time.

Suggested Answer: A

Community vote distribution

C (100%)

🗳️ 👤 **4555894** 6 months ago

Selected Answer: C

Create an SQS dead-letter queue. Modify the existing queue by including a re-drive policy that sets the `MaximumReceives` setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue. Instruct the scientists to use the dead-letter queue to review the data that is not valid. Reprocess this data at a later time.

upvoted 2 times

🗳️ 👤 **zijo** 6 months ago

Answer is A. Lambda function is required for automated fixing of the invalid message data and hence A is the right choice here.

upvoted 1 times

🗳️ 👤 **dzn** 6 months ago

This is not a good approach because it requires unifying the validation logic of the custom application and Lambda function, requires updating both the custom application and Lambda when data specifications change, and requires that the timing of those updates be the same from the SQS perspective, making the deployment process more complex and devops cost expensive. BTW, failed messages are reviewed by scientists, and there is no requirement that they be automatically fix by the program.

upvoted 1 times

🗳️ 👤 **dzn** 6 months, 3 weeks ago

Selected Answer: C

A Dead Letter Queue (DLQ) can be the destination queue for messages that cannot be successfully processed by other queues. DLQs are used to analyze why a message failed or to isolate problem messages.

upvoted 2 times

🗳️ 👤 **thanhv142** 7 months, 2 weeks ago

C is correct: Use dead letter queue and config maximum receives is the right way

upvoted 1 times

🗳️ 👤 **Bans** 8 months, 1 week ago

definitely C

upvoted 1 times

🗨️ 👤 **harithzainudin** 9 months ago

Selected Answer: C

This is DLQ use case. So, its 100% C
upvoted 3 times

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

Selected Answer: C

everyone votes C but answer seems as A. which one correct? should we trust to voters or examtopic? :D
upvoted 4 times

🗨️ 👤 **xhi158** 10 months, 2 weeks ago

The correct answer is C . This is a use case for Dead Letter Queue
upvoted 2 times

🗨️ 👤 **bugincloud** 11 months, 3 weeks ago

Selected Answer: C

classic DLQ usecase
upvoted 1 times

🗨️ 👤 **Skshitiz** 1 year ago

Selected Answer: C

C - DLQ
upvoted 1 times

🗨️ 👤 **FEEREWMWKA** 1 year ago

C - DLQ
upvoted 1 times

🗨️ 👤 **andriit** 1 year, 1 month ago

DevOps is about automation! Variant C says: "Instruct scientists... " :D
So variant A is the best among other
upvoted 1 times

🗨️ 👤 **Just_Ninja** 1 year, 1 month ago

Selected Answer: C

DLQ is the right solution. SQS is one to one! So Lambda make no sense.
upvoted 4 times

🗨️ 👤 **habros** 1 year, 2 months ago

Selected Answer: C

Always do with DLQ for failed deliveries. C all the way
upvoted 1 times

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

Selected Answer: C

C answer with DLQ is a right solution.
upvoted 1 times

🗨️ 👤 **bcx** 1 year, 3 months ago

Selected Answer: C

The perfect case for a dead-letter queue
upvoted 1 times

🗨️ 👤 **Akaza** 1 year, 3 months ago

Selected Answer: C


SQS DLQ needed
upvoted 1 times

A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application. Which solution ensures resources are deployed in accordance with company policy?

- A. Create AWS Trusted Advisor checks to find and remediate unapproved CloudFormation StackSets.
- B. Create a Cloud Formation drift detection operation to find and remediate unapproved CloudFormation StackSets.
- C. Create CloudFormation StackSets with approved CloudFormation templates.
- D. Create AWS Service Catalog products with approved CloudFormation templates.

Suggested Answer: D

Community vote distribution



harithzainudin Highly Voted 1 year, 3 months ago

Selected Answer: D

100% D.

AWS Service Catalog lets you centrally manage your cloud resources to achieve governance at scale of your infrastructure as code (IaC) templates, written in CloudFormation or Terraform configurations. With AWS Service Catalog, you can meet your compliance requirements while making sure your customers can quickly deploy the cloud resources they need.

<https://aws.amazon.com/servicecatalog/>

all other service in other answer is not related.

upvoted 9 times

flaacko Most Recent 6 months, 3 weeks ago

Service Catalogue is a like a internal marketplace for an organization in that accounts in that organization are limited to using only the resources describe in the product catalogue. For the use case described the best choice is using Service Catalogue.

upvoted 1 times

jamesf 7 months, 1 week ago

Selected Answer: D

keywords: strict tagging, resource requirements a, limit the deployment

AWS Service Catalog

upvoted 1 times

shammous 7 months, 2 weeks ago

D would be a better option, especially for developers, to abstract configuring the CloudFormation StackSets when launching applications with diverse versions. In AWS Service Catalog, they would just pick up the version and deploy. Everything would be set for them in the background including the CloudFormation StackSet with the version parameter and the tagging enforcement.

upvoted 1 times

Gomer 9 months, 1 week ago

Selected Answer: D

I'd argue that the correct answer is to use Service Catalog and StackSets. Option "D:" doesn't preclude using StackSets, it just doesn't mention it as part of the solution. ServiceCatalog is the formal method to distribute standard solutions (such as CloudFormation StackSets)

upvoted 1 times

Gomer 9 months, 1 week ago

"AWS Service Catalog enables you to launch a product in one or more accounts and AWS Regions. To do this, administrators must apply a stack set constraint to the product with the accounts and Regions, where it can launch as a stack set."

<https://docs.aws.amazon.com/servicecatalog/latest/userguide/launch-stacksets.html>

upvoted 1 times

stoy123 11 months, 2 weeks ago

"A provisioned Service Catalog product is an AWS CloudFormation stack"

Really confusing. I go with D...

upvoted 1 times

  **vn_thanhtung** 10 months, 1 week ago

https://docs.aws.amazon.com/servicecatalog/latest/adminguide/catalogs_constraints_template-constraints.html Please check topic this , correct answer is C

upvoted 1 times

  **zijo** 1 year ago

AWS Service Catalog can be used to deploy resources to two regions (or even more) with the help of AWS CloudFormation StackSets. So Answer is C

upvoted 3 times

  **Shasha1** 1 year ago

Answer C

If rules are applied across multiple accounts, the StackSets feature is more suitable. The service catalog is used for provisioning new accounts under the AWS control tower.

upvoted 3 times

  **Cert1Magic2** 1 year ago

Selected Answer: D

It's D

upvoted 2 times

  **dzn** 1 year ago

Selected Answer: C

Service Catalog cannot ensure that the same application can be deployed to multiple regions.

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago



D is correct: Catalog product impose strict requirements for app deployment. If using stacksets, devs can deploy app to everywhere without any restrictions

upvoted 1 times

  **thanhnv142** 1 year ago

Correction: should be C, not D. The question mentions <deployment to two Regions>. Only stacksets can do this. Even AWS Service Catalog products. It cannot be used cross-region unless it is deployed by stackset

upvoted 2 times

  **EricFu** 1 year, 1 month ago

To restrict regions or accounts where catalog products can be deployed, refer to AWS Service Catalog Stack Set Constraints (<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/constraints-stackset.html>),

upvoted 1 times

  **robertohyena** 1 year, 3 months ago

Selected Answer: D

I will go with D.

Reference here:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/manage-aws-service-catalog-products-in-multiple-aws-accounts-and-aws-regions.html>

upvoted 2 times

  **2pk** 1 year, 4 months ago

Selected Answer: D

While stacksets and catalog can use in this case catalog can restrict the strict policy than stacksets.

The company has strict tagging and resource requirements.. So it's product catalog. In stacksets developers can modify the resources since they can get hold the template.

upvoted 4 times

  **rlf** 1 year, 4 months ago

Answer is D.

"Developers will need to deploy" multiple versions of the same application. So service catalog products will be best for developers.

upvoted 3 times

  **BaburTurk** 1 year, 4 months ago

Selected Answer: C

The correct answer is: C. Create CloudFormation StackSets with approved CloudFormation templates.

CloudFormation StackSets enable you to deploy stacks across multiple accounts and Regions using a single template. This allows you to enforce company policy by only allowing developers to use approved templates.

AWS Service Catalog products can be used to launch approved CloudFormation templates, but they do not enforce the use of approved templates.

upvoted 2 times

  **Dushank** 1 year, 5 months ago

Selected Answer: C

The best solution to ensure that resources are deployed in accordance with company policy is to create CloudFormation StackSets with approved CloudFormation templates.

CloudFormation StackSets allow you to create and manage stacks across multiple AWS accounts and Regions. You can specify the template to use when creating a StackSet, as well as any parameters and capabilities that the template requires.

By using approved CloudFormation templates, you can ensure that all resources deployed by the StackSet meet your company's tagging and resource requirements. You can also use StackSets to limit the deployment to two Regions.

The other options are not as effective:

upvoted 2 times

A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.

Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

- A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zones. Update the route tables.
- B. Create additional EC2 instances spanning multiple Availability Zones. Add an Application Load Balancer to split the load between them.
- C. Configure an Application Load Balancer in front of the EC2 instance. Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure.
- D. Replace the NAT instance with a NAT gateway in each Availability Zone. Update the route tables.
- E. Replace the NAT instance with a NAT gateway that spans multiple Availability Zones. Update the route tables.

Suggested Answer: BD

Community vote distribution

BD (84%)

BE (16%)

  **Karamen** Highly Voted 1 year, 4 months ago

B&D

NAT Gateway does not span multiple AZ. you must create foreach AZ for HA

upvoted 9 times

  **HugoFM** Highly Voted 1 year, 3 months ago

BD



E is incorrect see NAT gateway basics in <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

upvoted 9 times

  **kaushald** 1 year ago

Quoting the above link: "If you have resources in multiple Availability Zones and they share one NAT gateway, and if the NAT gateway's Availability Zone is down, resources in the other Availability Zones lose internet access. To improve resiliency, create a NAT gateway in each Availability Zone, and configure your routing to ensure that resources use the NAT gateway in the same Availability Zone."

upvoted 3 times

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

Selected Answer: BD

<https://aws.amazon.com/blogs/networking-and-content-delivery/using-nat-gateways-with-multiple-amazon-vpcs-at-scale/>

NAT Gateways within an AZ are automatically implemented with redundancy. However, while Amazon VPCs can span multiple AZs, each NAT Gateway operates within a single AZ. If the NAT Gateway fails, then connections with resources using that NAT Gateway also fail. Therefore, it's recommended to deploy one NAT Gateway in each AZ and routing traffic locally within the same AZ.

upvoted 2 times

  **zijo** 1 year ago

Both NAT Gateway and NAT instance are regional resources. But NAT Gateway offers automatic deployment across Availability Zones, you might need to manually configure redundancy across Availability Zones for NAT Instances.

upvoted 1 times


  **thanhnv142** 1 year, 1 month ago

B and D are correct: We need to span EC2 to multiple avai zones and replace nat instance with nat gateway in each zone

B: span EC2 to multiple avai zones

D: replace nat instance with nat gateway

upvoted 2 times

  **Bans** 1 year, 2 months ago

Answer is B and D

upvoted 1 times

  **harithzainudin** 1 year, 3 months ago

Selected Answer: BD

Answer is B and D,

NAT gateways are regional services and do not span across Availability Zones. So, E is completely wrong.

upvoted 5 times

🗨️ **zolphar_z** 1 year, 3 months ago

Selected Answer: BD

NAT Gateway can't span in multiple regions, only in one subnet, I just tried it using the AWS Console

upvoted 6 times

🗨️ **harithzainudin** 1 year, 3 months ago

yes you're correct! I can confirm this. So BD is the correct answer

upvoted 1 times

🗨️ **DaddyDee** 1 year, 5 months ago

B&D <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-comparison.html>

upvoted 1 times

🗨️ **bugincloud** 1 year, 5 months ago

Selected Answer: BD

B & D is correct

NAT GW does not span across AZ, And has to be created in multi AZ for HA.

upvoted 3 times

🗨️ **RVivek** 1 year, 6 months ago

Selected Answer: BD

E is wrong. Natgateway cannot multiple subnets/zones

upvoted 4 times

🗨️ **cocegas** 1 year, 6 months ago

Selected Answer: BD

BD correct.

E: incorrect because NAT Gateway does not span multi AZ, you need to deploy it to different AZs. Not like LB that spans multiAZ automatically.

upvoted 5 times

🗨️ **FEEREWMWKA** 1 year, 6 months ago

Defo BD - Cannot be E as Nat Gateways sit in one subnet

upvoted 2 times

🗨️ **lakescix** 1 year, 6 months ago

B,D.

E is wrong because NAT Gateway is deployed to a single public subnet (Cannot span multiple AZs)

upvoted 3 times

🗨️ **mamila** 1 year, 7 months ago

Selected Answer: BD

E is wrong, NAT Gateway is a zonal resource.

upvoted 4 times

🗨️ **Certified101** 1 year, 7 months ago

Selected Answer: BD

A NAT Gateway is spun up in a single subnet that lives in a AZ. So you cannot build a NAT GW that spans multiple AZ's. You will need to build a NAT GW in EACH AZ to succeed. BD are the correct answers.

upvoted 2 times

🗨️ **Kiroo** 1 year, 8 months ago

Selected Answer: BE

Being honest DE are really similar

But BE looks more correct due to use the same language

upvoted 4 times

🗨️ **flaacko** 6 months, 3 weeks ago

B is correct because it says you should create instances in multiple AZs and then set up a load balancer to split traffic between them. E is wrong because NAT gateways cannot span availability zones.

upvoted 1 times

A DevOps engineer is building a multistage pipeline with AWS CodePipeline to build, verify, stage, test, and deploy an application. A manual approval stage is required between the test stage and the deploy stage. The development team uses a custom chat tool with webhook support that requires near-real-time notifications.

How should the DevOps engineer configure status updates for pipeline activity and approval requests to post to the chat tool?

- A. Create an Amazon CloudWatch Logs subscription that filters on CodePipeline Pipeline Execution State Change. Publish subscription events to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the chat webhook URL to the SNS topic, and complete the subscription validation.
- B. Create an AWS Lambda function that is invoked by AWS CloudTrail events. When a CodePipeline Pipeline Execution State Change event is detected, send the event details to the chat webhook URL.
- C. Create an Amazon EventBridge rule that filters on CodePipeline Pipeline Execution State Change. Publish the events to an Amazon Simple Notification Service (Amazon SNS) topic. Create an AWS Lambda function that sends event details to the chat webhook URL. Subscribe the function to the SNS topic.
- D. Modify the pipeline code to send the event details to the chat webhook URL at the end of each stage. Parameterize the URL so that each pipeline can send to a different URL based on the pipeline environment.

Suggested Answer: C

Community vote distribution

C (100%)

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

Selected Answer: C

The DevOps engineer should configure status updates for pipeline activity and approval requests to post to the chat tool by creating an Amazon EventBridge rule that filters on CodePipeline Pipeline Execution State Change. The events should be published to an Amazon Simple Notification Service (Amazon SNS) topic, and an AWS Lambda function should be created to send event details to the chat webhook URL. The function should be subscribed to the SNS topic. Option C is the correct answer.

Option A is incorrect because it suggests using CloudWatch Logs instead of EventBridge, which is not the optimal solution for this use case. Option B is incorrect because it suggests using CloudTrail instead of CodePipeline events, which is not relevant. Option D is incorrect because modifying the pipeline code is not necessary and adds unnecessary complexity.

upvoted 14 times

🗳️ 👤 **YucelFuat** Most Recent 6 months, 1 week ago

Selected Answer: C

Exam Tip : If you see that question is related to an Event or Action --> EventBridge

upvoted 4 times

🗳️ 👤 **dzr** 1 year ago

Selected Answer: C

API calls related to AWS CodePipeline are logged by CloudTrail. However, changes to the execution state of CodePipeline are events, not API calls. These events can be captured via Amazon EventBridge.

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

C is correct: Use lambda to send event detail to the chat webhook url. Subscribe lambda to SNS topic

A: Log subscription filter is for logging, not event

B: Should not use lambda with cloudtrail events

D: no need to modify pipeline code to send event at the end of each stage

upvoted 2 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

cloudtrail event cannot trigger lambda

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year ago

A: no way to collect log from code pipeline


upvoted 1 times

  **madperro** 1 year, 8 months ago

Selected Answer: C

C makes most sene.

upvoted 1 times

  **ele** 1 year, 10 months ago

Selected Answer: C

C right

upvoted 1 times

  **alce2020** 1 year, 10 months ago

Selected Answer: C

C it is

upvoted 2 times


A company's application development team uses Linux-based Amazon EC2 instances as bastion hosts. Inbound SSH access to the bastion hosts is restricted to specific IP addresses, as defined in the associated security groups. The company's security team wants to receive a notification if the security group rules are modified to allow SSH access from any IP address.

What should a DevOps engineer do to meet this requirement?

- A. Create an Amazon EventBridge rule with a source of `aws.cloudtrail` and the event name `AuthorizeSecurityGroupIngress`. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- B. Enable Amazon GuardDuty and check the findings for security groups in AWS Security Hub. Configure an Amazon EventBridge rule with a custom pattern that matches GuardDuty events with an output of `NON_COMPLIANT`. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- C. Create an AWS Config rule by using the `restricted-ssh` managed rule to check whether security groups disallow unrestricted incoming SSH traffic. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- D. Enable Amazon Inspector. Include the Common Vulnerabilities and Exposures-1.1 rules package to check the security groups that are associated with the bastion hosts. Configure Amazon Inspector to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Suggested Answer: C

Community vote distribution



ixdb Highly Voted 1 year, 7 months ago

A is right.

The Config rule `restricted-ssh` will not check the ingress rule that use the CIDR other than `0.0.0.0/0` and not notify anyone.

upvoted 18 times

csG13 1 year, 2 months ago

A would send a notification for ANY change in the security group. The question clearly states that wants only when `0.0.0.0/0` is allowed.

Therefore, should be C.

upvoted 10 times

hoazgazh 11 months ago

"a notification if the security group rules are modified to allow SSH access from any IP address"

from any IP address => so A is correct, any change in SG should send noti

upvoted 1 times

MarDog Highly Voted 1 year, 8 months ago

Selected Answer: A

I'm going to have to go with A on this one:

<https://aws.plainenglish.io/detecting-modifications-to-aws-ec2-security-groups-2ef8989a3350>

<https://repost.aws/knowledge-center/monitor-security-group-changes-ec2>

upvoted 8 times

teo2157 Most Recent 1 month, 2 weeks ago

Selected Answer: C

The key point here is "allow SSH access from any IP address" which is exactly "the `restricted-ssh` managed rule", said that, it's C

upvoted 2 times

teo2157 3 months, 2 weeks ago

Selected Answer: C

Very, very, very hard question. I think the key point here is the ANY, based on that, it's C

upvoted 1 times

steli0 3 months, 2 weeks ago

Selected Answer: C

A would be right if the "ANY" word describing all IPs (`0.0.0.0/0`) wasn't there. CloudTrail will notify you for any SG rule change.

upvoted 2 times

🗳️ 👤 **BrusingWayne** 3 months, 3 weeks ago

Option C (Incorrect):

AWS Config rules are good for ongoing compliance checks, but they don't provide real-time notifications for changes.

Config rules run periodically, which could result in a delay between the change and the notification.

The automatic remediation aspect is not required in this scenario and could potentially interfere with legitimate changes.

Hence, it is Option A.

upvoted 1 times

🗳️ 👤 **Impromptu** 3 months, 3 weeks ago

Selected Answer: C

A: Would send a message to SNS for every change, so not only SSH but all other ports/services. This would be too much.

I do get the other comments that C would only notify for 0.0.0.0/0 but I think that is what the question is trying to state with "any IP".

upvoted 3 times

🗳️ 👤 **anuvindh** 5 months ago

C is the answer : <https://docs.aws.amazon.com/config/latest/developerguide/notifications-for-AWS-Config.html>

upvoted 3 times

🗳️ 👤 **jamesf** 7 months, 1 week ago

Selected Answer: C

keywords: Inbound SSH access

C restricted for SSH port (22) only from ANY address

upvoted 3 times

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

A! "AWS Config provides rules such as restricted-ssh that can be used to detect Security Groups that have SSH access open for any IP".

upvoted 1 times

🗳️ 👤 **TioChico** 8 months, 1 week ago

Selected Answer: A

A is right

upvoted 1 times

🗳️ 👤 **Sodev** 9 months, 3 weeks ago

Selected Answer: A

I think keyword for C must be "ALL".

ANY means when new IP is added to security group, so SNS will be triggered

upvoted 2 times

🗳️ 👤 **flaacko** 6 months, 3 weeks ago

In the context of AWS when you see ANY IP address, it is probably referring to the 0.0.0.0/0 CIDR block which allows traffic from all or any IP address from the internet. When you use the restricted-ssh managed rule, Security Groups will be labelled as NON_COMPLIANT when they allow unrestricted SSH traffic from anywhere or any IP address (0.0.0.0/0).

upvoted 1 times

🗳️ 👤 **liuyomz** 10 months ago

Selected Answer: C

C makes way more sense from the way AWS wants us to do it

upvoted 2 times

🗳️ 👤 **seetpt** 10 months, 2 weeks ago

Selected Answer: C

i vote for c

upvoted 2 times

🗳️ 👤 **c3518fc** 11 months ago

Selected Answer: A

A. This is the correct solution because it leverages Amazon EventBridge to monitor for changes to the security group rules, specifically the AuthorizeSecurityGroupIngress event, which indicates that the security group rules have been modified to allow SSH access from any IP address.

By creating an EventBridge rule with the appropriate event pattern and defining an Amazon SNS topic as the target, the DevOps engineer can ensure that the security team receives a notification whenever the security group rules are modified in an undesirable way.

upvoted 2 times

🗨️ 👤 **zijo** 11 months, 3 weeks ago

Answer is C

The restricted-ssh managed rule in AWS Config helps ensure your bastion host security groups are locked down for SSH access. It specifically checks if incoming SSH traffic is accessible for the security groups.

The rule is considered COMPLIANT if:

SSH access is not open to the public (meaning the rule doesn't find a security group allowing 0.0.0.0/0 for port 22).

SSH access is restricted to specific IP addresses or security groups using CIDR notation (e.g., 10.0.0.0/16).

If the rule detects a security group allowing SSH access from anywhere (0.0.0.0/0), it triggers a NON_COMPLIANT status.

upvoted 2 times

🗨️ 👤 **c3518fc** 11 months ago

Yeah, but has nothing to do with anyone changing it. A is your answer because it detects changes and sends out an email notification

upvoted 1 times

🗨️ 👤 **Cervus18** 11 months, 3 weeks ago

Selected Answer: C

restricted-ssh : The rule is COMPLIANT if the IP addresses of the incoming SSH traffic in the security groups are restricted (CIDR other than 0.0.0.0/0 or ::/0). Otherwise, NON_COMPLIANT.

<https://docs.aws.amazon.com/config/latest/developerguide/restricted-ssh.html>

That addresses exactly the requirement !

upvoted 3 times


A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.

Which actions should be taken to accomplish this? (Choose two.)

- A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
- B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.
- C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.
- D. Modify the on-premises application to send log information back to API Gateway with each request.
- E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

Suggested Answer: AC

Community vote distribution



madperro Highly Voted 1 year, 8 months ago

Selected Answer: AC

AC is using standard parts of the solution.

upvoted 8 times

thanhnv142 Highly Voted 1 year, 1 month ago

A and C: use cloudwatch log agent to collect app log and use AWS X-ray to collect information about requests (traces).

B is incorrect because modifying app to send message directly to X-RAY introduces more latency to the app. Use X-RAY daemon to do that task is a better idea

upvoted 5 times

thanhnv142 1 year, 1 month ago

- <API Gateway latency metrics> means dev team have collect APM. They need to collect app log as well, which indicates option A.

D and E: modifying the app introduces latencies

upvoted 2 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: AC

AC is less impact to Application Latencies.

Keywords: without additional latencies, cloudwatch

B will provide more latencies

DE will require modify the app and give more latencies.

upvoted 1 times

TEC1 10 months, 2 weeks ago

Selected Answer: AC

The X-Ray daemon batches and uploads the data in the background, which helps to avoid introducing additional latency.

upvoted 3 times

yorkicurke 1 year, 3 months ago

Selected Answer: AC

the reason i am not so sure about is that API Gateway have built-in integration with X-Ray. This means that they automatically send trace data to X-Ray without needing a separate X-Ray daemon. and i dont think we have the option of installing one or using one, unless someone shows me the official link.

upvoted 3 times

Bassel 1 year, 9 months ago



Selected Answer: BE

Installing the CloudWatch agent server-side (option A) is not directly related to collecting latency data from API Gateway. The CloudWatch agent is typically used to collect and monitor system-level metrics from the server itself.

Enabling AWS X-Ray tracing in API Gateway and using the X-Ray daemon (option C) is not necessary in this scenario. The X-Ray daemon is primarily used when you have applications running on EC2 instances or on-premises servers that need to send trace data to X-Ray.

Modifying the on-premises application to send log information back to API Gateway with each request (option D) is not an optimal solution for collecting latency data. It may introduce additional latency and overhead to the API requests and could be challenging to implement efficiently and accurately.

upvoted 3 times

  **rhinozD** 1 year, 9 months ago

Do you think that doing B or E doesn't bring any latency?

I think C is necessary because you could trace the performance of the application.



And even the team can look into app logs on its server, but sending logs to Cloudwatch logs and then making a further investigation with AWS tools is not too bad.

upvoted 2 times

  **EricZhang** 1 year, 9 months ago



Why A? The team still can check logs without uploading to CloudWatch? I'd prefer E over A.

upvoted 2 times

  **NivNZ** 1 year, 7 months ago

I thought the same but E might cause additional latency which is NOT what we want.



upvoted 3 times

  **ele** 1 year, 10 months ago

Selected Answer: AC

AC less impact on app

upvoted 4 times

  **alce2020** 1 year, 10 months ago

Selected Answer: AC

A, C, correct

upvoted 2 times

A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.

Which solution will accomplish this?

- A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoints. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.
- B. Create an Aurora custom endpoint to point to the primary database instance. Configure the application to use this endpoint. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
- C. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instance. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
- D. Store the Aurora endpoint in AWS Systems Manager Parameter Store. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Store. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Suggested Answer: D

Community vote distribution

D (100%)

 **haazybanj**  1 year, 10 months ago

Selected Answer: D

D is the correct answer.

Explanation:

To automate the promotion of a read replica to the primary instance in the event of a failure, we need to detect the failure and then invoke an AWS Lambda function to promote the replica instance. This can be achieved using Amazon EventBridge.

Option A is incorrect because using a CNAME with health checks doesn't provide an automated way to promote the read replica. Additionally, subscribing an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail doesn't help to promote the replica.

Option B is incorrect because a custom endpoint is not required to promote the read replica. Additionally, using AWS CloudTrail to run an AWS Lambda function to promote the replica instance doesn't provide an automated way to update the application endpoint to point to the newly promoted instance.

upvoted 8 times

 **n1w**  1 year, 4 months ago

doesn't failover happen automatically in aurora?

upvoted 7 times

 **VrilianVirgil** 1 year ago

Aurora supports automated failover for a single cluster. [Be it a global Aurora cluster or a multi AZ/region deployment]

In this case it's implied that the read-replica is not part of the cluster.

that's my best guess.

upvoted 2 times

 **jamesf**  7 months, 2 weeks ago

Selected Answer: D

D is correct.

Option B is wrong as AWS CloudTrail to run an AWS Lambda function to promote the replica instance doesn't provide an automated way.

upvoted 1 times

🗨️ **jamesf** 7 months, 1 week ago

Option B is wrong also due to:

- Custom Endpoint Management: Extra complexity in managing and updating endpoints dynamically.
- Lag in Promotion: Possible delays due to CloudTrail event delivery and Lambda invocation.
- Reliance on CloudTrail: Lag in event processing can cause potential downtime or data inconsistency.

upvoted 1 times

🗨️ **hkh2** 8 months ago

Correct answer is B

Here is why.

Previously, you might have used the CNAME mechanism to set up Domain Name Service (DNS) aliases from your own domain to achieve similar results. By using custom endpoints, you can avoid updating CNAME records when your cluster grows or shrinks. Custom endpoints also mean that you can use encrypted Transport Layer Security/Secure Sockets Layer (TLS/SSL) connections.

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.Endpoints.html#Aurora.Endpoints.Custom>

upvoted 1 times

🗨️ **thanhnv142** 1 year, 1 month ago

A is correct: Using Amazon Route 53 CNAME with health checks is the way for failover recommended by AWS:

<https://aws.amazon.com/blogs/database/cross-region-disaster-recovery-using-amazon-aurora-global-database-for-amazon-aurora-postgresql/>

upvoted 1 times

🗨️ **Ffida** 1 year, 5 months ago

option D is not either providing seamless solution, in option D application needed to be reload and that will cause downtime.

upvoted 1 times

🗨️ **madperro** 1 year, 8 months ago

Selected Answer: D

D make most sense.

upvoted 2 times

🗨️ **haazybanj** 1 year, 10 months ago

Selected Answer: D

Option D is the correct solution

Option C is incorrect because modifying the AWS CloudFormation template requires manual intervention and cannot be automated. Additionally, creating an Amazon CloudWatch alarm to invoke the Lambda function after the failure event occurs doesn't provide an automated way to promote the replica instance.

Therefore, Option D is the correct solution.

upvoted 3 times

🗨️ **haazybanj** 1 year, 10 months ago

Selected Answer: D

D is the answer

upvoted 2 times

🗨️ **mgonblan** 1 year, 10 months ago

D: Reference:<https://aws.amazon.com/es/blogs/database/cross-region-cross-account-disaster-recovery-using-amazon-aurora-global-database/>

upvoted 2 times

🗨️ **alce2020** 1 year, 11 months ago

Selected Answer: D

D it is


upvoted 2 times

A company hosts its staging website using an Amazon EC2 instance backed with Amazon EBS storage. The company wants to recover quickly with minimal data losses in the event of network connectivity issues or power failures on the EC2 instance. Which solution will meet these requirements?

- A. Add the instance to an EC2 Auto Scaling group with the minimum, maximum, and desired capacity set to 1.
- B. Add the instance to an EC2 Auto Scaling group with a lifecycle hook to detach the EBS volume when the EC2 instance shuts down or terminates.
- C. Create an Amazon CloudWatch alarm for the StatusCheckFailed System metric and select the EC2 action to recover the instance.
- D. Create an Amazon CloudWatch alarm for the StatusCheckFailed Instance metric and select the EC2 action to reboot the instance.

Suggested Answer: A

Community vote distribution



madperro Highly Voted 1 year, 8 months ago

Selected Answer: C

C is the right answer.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

upvoted 11 times

tartarus23 Highly Voted 1 year, 8 months ago

Selected Answer: C

Explanation:

Amazon CloudWatch provides system-wide visibility into resource utilization, application performance, and operational health. If a system status check fails, this implies there's a problem with the underlying EC2 system that may require AWS involvement to repair. The "Recover this instance" action for the system status check automatically recovers the instance if it becomes impaired due to an underlying issue.

upvoted 10 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: C

C

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

In the event that AWS determines an instance is unavailable due to an underlying hardware issue, there are two mechanisms that you can configure for instance resiliency which can restore availability—simplified automatic recovery and Amazon CloudWatch action based recovery. This process is called instance recovery.

The following are examples of underlying hardware issues that might require instance recovery:

- Loss of network connectivity
- Loss of system power
- Software issues on the physical host
- Hardware issues on the physical host that impact network reachability

upvoted 1 times

c3518fc 11 months ago

Selected Answer: C

C. This is the correct solution. By creating a CloudWatch alarm for the StatusCheckFailed System metric and configuring the alarm to trigger the "Recover this instance" action, the EC2 instance will be automatically recovered in the event of a system failure or power outage. This ensures the instance can be quickly recovered with minimal data loss, as the EBS volume remains attached during the recovery process.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

upvoted 1 times

4555894 1 year ago

Selected Answer: C

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

C is correct: recover is the right way

A and B are irrelevant

D: should not reboot in case of power failures

upvoted 1 times

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

Selected Answer: C

My Reason:

StatusCheckFailed_System: This check monitors the AWS systems on which your instance runs¹. For Example loss of network connectivity, loss of system power, software issues on the physical host, and hardware issues on the physical host that impact network reachability

StatusCheckFailed_Instance: This check monitors the software and network configuration of your individual instance. These checks detect problems that require your involvement to repair. If an instance status check fails, it typically means that there's an issue with the instance, such as a misconfigured network or a problem with the instance's file system.

upvoted 3 times

🗳️ 👤 **bakamon** 1 year, 8 months ago

Selected Answer: C

Correct Answer is C

upvoted 2 times

🗳️ 👤 **qan1257** 1 year, 9 months ago

Selected Answer: C

A is incorrect.

Simplified automatic recovery is not initiated for instances in an Auto Scaling group. If your instance is part of an Auto Scaling group with health checks enabled, then the instance is replaced when it becomes impaired.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

upvoted 4 times

🗳️ 👤 **ele** 1 year, 10 months ago

Selected Answer: C

C with recover action creates identical instance

upvoted 1 times

🗳️ 👤 **Zoe_zoe** 1 year, 10 months ago

C

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

There are only 2 ways to recover EC2 instances. Since this instance has EBS volumes, only the CloudWatch action based recovery is applicable.

upvoted 3 times

🗳️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: D

Option D is the correct solution

Option C is incorrect because modifying the AWS CloudFormation template requires manual intervention and cannot be automated. Additionally, creating an Amazon CloudWatch alarm to invoke the Lambda function after the failure event occurs doesn't provide an automated way to promote the replica instance.

Therefore, Option D is the correct solution.

upvoted 1 times

🗳️ 👤 **alce2020** 1 year, 11 months ago

C is correct

upvoted 1 times


🗳️ 👤 **jqso234** 1 year, 11 months ago

Selected Answer: A

option A is a better choice for this scenario because it ensures that there is always an EC2 instance running to serve the staging website, and the new instance will have the same configuration as the original instance, including the EBS volume, so there will be minimal data loss. Option C

may result in some data loss since a new EBS volume will be created, and it may take longer to recover the instance since the EC2 action to recover the instance will need to be triggered by the Amazon CloudWatch alarm.

upvoted 1 times

  **bcx** 1 year, 9 months ago

You would lose the contents on the EBS volume.

upvoted 1 times

  **ogwu2000** 1 year, 7 months ago

How would you launch a new instance if they is a power outage? So, C is correct as you will have to recover and hopefully quickly.

upvoted 1 times



  **Ffida2214** 1 year, 2 months ago

If there is power outage than it is considered as instance failure and cloudwatch alarm can recover from system failure but can't recover from instance failure.

<https://repost.aws/knowledge-center/automatic-recovery-ec2-cloudwatch>

I believe option A is quickest recovery, and if Data needed to be backup then option B

upvoted 1 times

  **Dimidrol** 1 year, 11 months ago

Selected Answer: C

C for me


upvoted 3 times

A company wants to use AWS development tools to replace its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates file permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services. Which solution will meet these requirements?

- A. Use AWS CodeBuild to test the application. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, and deregister and register instances with the ALB. Use the appspec.yml file to update file permissions without a custom script.
- B. Use AWS CodePipeline to move the application from the AWS CodeCommit repository to AWS CodeDeploy. Use CodeDeploy's deployment group to test the application, unregister and re-register instances with the ALB and restart services. Use the appspec.yml file to update file permissions without a custom script.
- C. Use AWS CodePipeline to move the application source code from the AWS CodeCommit repository to AWS CodeDeploy. Use CodeDeploy to test the application. Use CodeDeploy's appspec.yml file to restart services and update permissions without a custom script. Use AWS CodeBuild to unregister and re-register instances with the ALB.
- D. Use AWS CodePipeline to trigger AWS CodeBuild to test the application. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services. Unregister and re-register the instances in the AWS CodeDeploy deployment group with the ALB. Update the appspec.yml file to update file permissions without a custom script.

Suggested Answer: D

Community vote distribution



madperro Highly Voted 1 year, 8 months ago

Selected Answer: D

D is better than A. You need to include CodePipeline to move execution from CodeBuild to CodeDeploy.
upvoted 9 times

haazybanj Highly Voted 1 year, 10 months ago

Selected Answer: D

Option D is also a viable solution. It suggests using AWS CodePipeline to trigger AWS CodeBuild to test the application, and then use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, unregister and re-register instances with the ALB, and update file permissions. This approach also covers all the deployment functionality required by the company
upvoted 5 times

rk0509 Most Recent 6 months, 4 weeks ago

Selected Answer: B

Answer is B. company want to replace its bash deployment scripts so option D is not suitable
upvoted 1 times

SabeloM 2 months, 3 weeks ago

Options D is suitable since it includes "Update the appspec.yml file to update file permissions without a custom script".
upvoted 1 times

rk0509 6 months, 4 weeks ago

Answer is B. company want to replace its bash deployment scripts so option D is not suitable
upvoted 2 times

jamesf 7 months, 2 weeks ago

Selected Answer: D

Should be D

CodePipeline - execute from CodeBuild to CodeDeploy

CodeBuild - test the application

CodeDeploy - deploy app, restart services, Unregister and re-register instance

Not Option A: not using CodePipeline

Not Option BC: using CodeCommit repo, not relevant with question.

upvoted 3 times

🗨️ 👤 **zijo** 11 months, 3 weeks ago

codebuild to test not codedeploy D is correct

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

D: is correct: need codepipeline for a seamless deployment. Need codebuild to test and codedeploy to deploy the app on EC2

A: no mention of codepipeline

B and C both mention AWS CodeCommit repository, which is irrelevant

upvoted 4 times

🗨️ 👤 **thanhnv142** 1 year ago

A: <deregister and register instances with the ALB>: we need to unregister, not deregister it

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

B and C: The question doesnt mention the need for a source code repository.

B: move the application from the AWS CodeCommit repository to AWS CodeDeploy -> Cannot do this, codecommit does not store apps, only code

C: move the application source code from the AWS CodeCommit repository to AWS CodeDeploy -> cannot do this, code deploy does not store code

upvoted 1 times

🗨️ 👤 **tartarus23** 1 year, 8 months ago

Selected Answer: A

Explanation:

AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy, which is perfect for unit testing the application.

AWS CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances. You can specify scripts to be run at set points during a deployment lifecycle, such as deregistering and registering instances with a load balancer, stopping and starting services, or changing file permissions, by defining them in the appspec.yml file.

upvoted 2 times

🗨️ 👤 **c3518fc** 11 months ago

but it says to deregister

upvoted 1 times

🗨️ 👤 **bakamon** 1 year, 8 months ago

Selected Answer: D

D is the correct bubamon

upvoted 2 times

🗨️ 👤 **Akaza** 1 year, 9 months ago

D for sure

upvoted 1 times

🗨️ 👤 **alce2020** 1 year, 11 months ago

Selected Answer: D

D it is


upvoted 2 times

A company runs an application with an Amazon EC2 and on-premises configuration. A DevOps engineer needs to standardize patching across both environments. Company policy dictates that patching only happens during non-business hours. Which combination of actions will meet these requirements? (Choose three.)

- A. Add the physical machines into AWS Systems Manager using Systems Manager Hybrid Activations.
- B. Attach an IAM role to the EC2 instances, allowing them to be managed by AWS Systems Manager.
- C. Create IAM access keys for the on-premises machines to interact with AWS Systems Manager.
- D. Run an AWS Systems Manager Automation document to patch the systems every hour
- E. Use Amazon EventBridge scheduled events to schedule a patch window.
- F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

Suggested Answer: ABF

Community vote distribution



thanhnv142 Highly Voted 1 year, 1 month ago

ABF are the right answers:

A: enable hybrid on AWS system manager

B: create IAM role for System manager to manage EC2 instances

F: use maintenance windows to schedule patching on non-business hours

C: incorrect because there is no IAM access keys for on-prem

D: should not run patching every hour

E: should not use Eventbridge because AWS has its own service to schedule patching

upvoted 9 times

DavidPham Highly Voted 1 year, 7 months ago

Selected Answer: ABF

ABF is correct

upvoted 5 times

79f3aa3 Most Recent 2 weeks, 2 days ago

Selected Answer: ABF

ABF is correct. But option 'B' should also include on-prem servers. Both EC2 and on-prem servers will use the hybrid activation role.

upvoted 1 times

spring21 2 months, 2 weeks ago

Selected Answer: ACF

To create IAM access keys for on-premises machines to interact with AWS Systems Manager, you need to: create a dedicated IAM user with the necessary permissions for Systems Manager actions, then generate access keys for that user and securely store them on the on-premises machine; ensure you follow best practices like rotating access keys regularly and using a secure method to distribute them.

upvoted 2 times

jamesf 7 months, 2 weeks ago

Selected Answer: ABF

ABF are correct

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

To configure non-EC2 machines for use with AWS Systems Manager in a hybrid and multicloud environment, you create a hybrid activation. Non-EC2 machine types supported as managed nodes include the following:

- Servers on your own premises (on-premises servers)
- AWS IoT Greengrass core devices
- AWS IoT and non-AWS edge devices
- Virtual machines (VMs), including VMs in other cloud environments

upvoted 2 times

🗨️ 👤 **HarryLy** 9 months, 1 week ago

ABF is correct

upvoted 1 times

🗨️ 👤 **Kiroo** 1 year, 7 months ago

Selected Answer: AF

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

AF are right but the letter B is wrong the role is for non EC2 instances

upvoted 2 times

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

Selected Answer: ABF

ABF is correct.

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

upvoted 4 times

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

Selected Answer: ABF

ABF is right

upvoted 3 times

🗨️ 👤 **alice2020** 1 year, 11 months ago

Selected Answer: ABF

ABF it is

upvoted 3 times

A company has chosen AWS to host a new application. The company needs to implement a multi-account strategy. A DevOps engineer creates a new AWS account and an organization in AWS Organizations. The DevOps engineer also creates the OU structure for the organization and sets up a landing zone by using AWS Control Tower.


The DevOps engineer must implement a solution that automatically deploys resources for new accounts that users create through AWS Control Tower Account Factory. When a user creates a new account, the solution must apply AWS CloudFormation templates and SCPs that are customized for the OU or the account to automatically deploy all the resources that are attached to the account. All the OUs are enrolled in AWS Control Tower.

Which solution will meet these requirements in the MOST automated way?

- A. Use AWS Service Catalog with AWS Control Tower. Create portfolios and products in AWS Service Catalog. Grant granular permissions to provision these resources. Deploy SCPs by using the AWS CLI and JSON documents.
- B. Deploy CloudFormation stack sets by using the required templates. Enable automatic deployment. Deploy stack instances to the required accounts. Deploy a CloudFormation stack set to the organization's management account to deploy SCPs.
- C. Create an Amazon EventBridge rule to detect the CreateManagedAccount event. Configure AWS Service Catalog as the target to deploy resources to any new accounts. Deploy SCPs by using the AWS CLI and JSON documents.
- D. Deploy the Customizations for AWS Control Tower (CfCT) solution. Use an AWS CodeCommit repository as the source. In the repository, create a custom package that includes the CloudFormation templates and the SCP JSON documents.

Suggested Answer: D

Community vote distribution



tartarus23 Highly Voted 1 year, 8 months ago

Selected Answer: D

The CfCT solution is designed for the exact purpose stated in the question. It extends the capabilities of AWS Control Tower by providing you with a way to automate resource provisioning and apply custom configurations across all AWS accounts created in the Control Tower environment. This enables the company to implement additional account customizations when new accounts are provisioned via the Control Tower Account Factory.

The CloudFormation templates and SCPs can be added to a CodeCommit repository and will be automatically deployed to new accounts when they are created. This provides a highly automated solution that does not require manual intervention to deploy resources and SCPs to new accounts.

upvoted 7 times

madperro Highly Voted 1 year, 8 months ago

Selected Answer: D

CfCT is designed for the purpose stated in the question. So D.

<https://docs.aws.amazon.com/controltower/latest/userguide/cfct-overview.html>

upvoted 6 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: D

D

<https://docs.aws.amazon.com/controltower/latest/userguide/cfct-overview.html>

Customizations for AWS Control Tower (CfCT) helps you customize your AWS Control Tower landing zone and stay aligned with AWS best practices. Customizations are implemented with AWS CloudFormation templates and service control policies (SCPs).

upvoted 1 times

jamesf 7 months, 2 weeks ago

keywords: "sets up a landing zone by using AWS Control Tower"

upvoted 1 times

Gomer 9 months, 1 week ago

Selected Answer: D

"This CfCT capability is integrated with AWS Control Tower lifecycle events, so that your resource deployments remain synchronized with your landing zone."

"For example, when a new account is created through account factory, all resources attached to the account are deployed automatically."

"You can deploy the custom templates and policies to individual accounts and organizational units (OUs) within your organization."

<https://docs.aws.amazon.com/controltower/latest/userguide/cfct-overview.html>

upvoted 1 times

🗲️ 👤 **thanhnv142** 1 year, 1 month ago

D is correct: Use CfCT is the correct solution: it utilizes both CloudFormation template and SCP

A and C: no mention of AWS CloudFormation

B: No mention of AWS control tower

upvoted 3 times

🗲️ 👤 **khchan123** 1 year, 1 month ago

Selected Answer: D

D. B is wrong because StackSets doesn't deploy stack instances to the organization management account.

upvoted 3 times

🗲️ 👤 **Bassel** 1 year, 9 months ago

Selected Answer: B

B. Deploying CloudFormation stack sets is the most automated way to deploy resources for new accounts created through AWS Control Tower Account Factory. With stack sets, you can define a CloudFormation template and deploy it to multiple accounts automatically. By enabling automatic deployment and deploying stack instances to the required accounts, you can ensure that the resources specified in the CloudFormation templates are automatically provisioned for each account. Additionally, by deploying a CloudFormation stack set to the organization's management account, you can deploy Service Control Policies (SCPs) across all accounts in the organization.

upvoted 2 times

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

Customizations for AWS Control Tower combines AWS Control Tower and other highly-available, trusted AWS services to help customers more quickly set up a secure, multi-account AWS environment using AWS best practices. You can easily add customizations to your AWS Control Tower landing zone using an AWS CloudFormation template and service control policies (SCPs). You can deploy the custom template and policies to individual accounts and organizational units (OUs) within your organization. It also integrates with AWS Control Tower lifecycle events to ensure that resource deployments stay in sync with your landing zone. For example, when a new account is created using the AWS Control Tower account factory, Customizations for AWS Control Tower ensures that all resources attached to the account's OUs will be automatically deployed.

upvoted 2 times

🗲️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: D

D is it

upvoted 2 times

🗲️ 👤 **alce2020** 1 year, 11 months ago

Selected Answer: D

D it is

upvoted 2 times

An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.

When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.

How should the company meet these requirements with the LEAST amount of application changes?

- A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.
- B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.
- C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.
- D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

Suggested Answer: C

Community vote distribution

C (100%)

 **Bassel** Highly Voted 1 year, 9 months ago

Selected Answer: C

C. Using Aurora with read replicas for the product catalog allows for a single product catalog across all regions. Aurora read replicas can be set up in different regions to provide low-latency access to the product catalog from each region. Additionally, by deploying additional local Aurora instances in each region for customer information and purchases, the company can comply with the requirement of keeping customer data and purchases in each region.

upvoted 9 times

 **dark4igi** Most Recent 1 week, 4 days ago

Selected Answer: C

C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.

It's very bad to use word "local" in answer about AWS

upvoted 1 times

 **jamesf** 7 months, 2 weeks ago

Selected Answer: C

C

keywords: "the LEAST amount of application changes"

upvoted 2 times

 **Sisanda_giiven** 1 year ago

Selected Answer: C

How should the company meet these requirements with the LEAST amount of application changes? Anything option with DynamoDB is out since the all the data is stored Aurora(relational database).

upvoted 3 times

 **thanhnv142** 1 year, 1 month ago

C is correct: data is kept in each region and one product catalog for all regions

A: Redshift is for data analysis, not for the need in the question

B: DynamoDB is primarily used for session data in a web app

D: Amazon DynamoDB global tables for the customer information is against the policy

upvoted 3 times

 **amrit1227** 1 year, 2 months ago

C is correct

upvoted 1 times

 **madperro** 1 year, 8 months ago

Selected Answer: C

C makes most sense and minimizes application changes.


upvoted 2 times

  **haazybanj** 1 year, 10 months ago

Selected Answer: C

The best solution to meet the company's requirements with the LEAST amount of application changes is to use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases. This will allow for a single product catalog across all regions, while still keeping customer information and purchases in each region for compliance purposes. Amazon Redshift is a data warehousing solution and is not appropriate for this use case. Amazon DynamoDB global tables may be used, but they require application changes to support them. Using local Aurora instances in each region for customer information and purchases could also work, but this would require more configuration and management than using Aurora with read replicas. Therefore, option C is the best solution.

upvoted 3 times

  **alce2020** 1 year, 11 months ago

Selected Answer: C

C is correct

upvoted 3 times

A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.

The API stack contains the following three tiers:

Amazon API Gateway -

AWS Lambda -

Amazon DynamoDB -


Which solution will meet the requirements?

- A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
- B. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
- C. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Region. Retrieve the data from a DynamoDB global table. Deploy a Lambda function to check the North America API health every 5 minutes. In the event of a failure, update Route 53 to point to the disaster recovery API.
- D. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routing. Configure the API to forward requests to the Lambda function in the Region nearest to the user. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

Suggested Answer: B

Community vote distribution

B (100%)

 **haazybanj** Highly Voted 1 year, 4 months ago


Selected Answer: B

B is the correct solution.

The requirement is to ensure both high reliability and fast response times for users located in North America and Europe. To meet this requirement, we can use Amazon Route 53 with latency-based routing to direct users to the closest API Gateway endpoint. Additionally, we can use health checks to monitor the health of each endpoint and direct traffic away from unhealthy endpoints.

To maintain high reliability, we can use AWS Lambda to handle the API requests. Since Lambda scales automatically, we don't need to worry about provisioning or maintaining infrastructure. We can also use DynamoDB as the database since it provides low latency access and automatic scaling.

upvoted 12 times

 **thanhnv142** Most Recent 7 months, 2 weeks ago


B is correct: using both latency-based routing and health checks ensures high reliability and fast response

A: only health check doesn't ensure fast response

C: All traffic would be routed to one location only (either NA or Europe if NA failed)


D: All traffic would be routed to one NA only. There would be no entry point which is near Europe users.

upvoted 3 times

 **amrit1227** 8 months, 2 weeks ago

B is correct

upvoted 1 times

 **z_inderjot** 8 months, 3 weeks ago

Selected Answer: B

B ,

Using latency based routing for better response time . Having api gateway in each region reduce request flight time. Lambda and Dynamo being a managed service scale automatically and having them in same region just reduce latency.

upvoted 2 times

🗲️ 👤 **Snape** 1 year, 1 month ago

Selected Answer: B

Reliability is different than resiliency, hence A and C are out as they are focussing on health checks which is required for the resiliency. DR again for the resiliency

upvoted 3 times

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

Selected Answer: B

B is the best solution.

upvoted 4 times

🗲️ 👤 **alice2020** 1 year, 4 months ago

Selected Answer: B

B is correct

upvoted 2 times

A rapidly growing company wants to scale for developer demand for AWS development environments. Development environments are created manually in the AWS Management Console. The networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.

To keep up with demand, the DevOps engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the development environments.

Which approach will meet these requirements and quickly provide consistent AWS environments for developers?

- A. Use `Fn::ImportValue` intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet values. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments needed. Use the `UpdateStackSet` command to update existing development environments.
- B. Use nested stacks to define common infrastructure components. To access the exported values, use `TemplateURL` to reference the networking team's template. To retrieve Virtual Private Cloud (VPC) and subnet values, use `Fn::ImportValue` intrinsic functions in the Parameters section of the root template. Use the `CreateChangeSet` and `ExecuteChangeSet` commands to update existing development environments.
- C. Use nested stacks to define common infrastructure components. Use `Fn::ImportValue` intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet values. Use the `CreateChangeSet` and `ExecuteChangeSet` commands to update existing development environments.
- D. Use `Fn::ImportValue` intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet values. Define the development resources in the order they need to be created in the CloudFormation nested stacks. Use the `CreateChangeSet` and `ExecuteChangeSet` commands to update existing development environments.

Suggested Answer: C

Community vote distribution

C (77%)

B (23%)

  **ipsingh** Highly Voted 1 year, 3 months ago

C is Correct.

B is WRONG because intrinsic functions can't be used in Parameter as per AWS documentation.

<https://repost.aws/knowledge-center/cloudformation-template-validation>

upvoted 14 times

  **aksliveswithaws** 12 months ago

You can use intrinsic functions only in specific parts of a template. Currently, you can use intrinsic functions in resource properties, outputs, metadata attributes, and update policy attributes

Refer

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference.html>

upvoted 2 times

  **seetpt** Most Recent 4 months, 1 week ago

Selected Answer: C

C seems right

upvoted 1 times

  **thanhnv142** 7 months, 1 week ago

C is correct: use nested stacks and `Fn::ImportValue` intrinsic functions with the resources of the nested stack

A: no mention of nested stack

B and D: `Fn::ImportValue` intrinsic function is used on child template to import values from parent template. So it should not be used on root template, which is the universal parent template of all other templates

upvoted 4 times

  **madperro** 1 year, 2 months ago

Selected Answer: C

C is the best answer. B is wrong as you need to use `Fn::ImportValue` in Resource section to import CFN template outputs.

upvoted 2 times

  **ducluanxutrieu** 1 year, 2 months ago

Selected Answer: C

I will go with C

upvoted 3 times

  **tartarus23** 1 year, 2 months ago

Selected Answer: B

B. Use nested stacks to define common infrastructure components. To access the exported values, use TemplateURL to reference the networking team's template. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root template. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.

Nested stacks allow you to modularize and reuse CloudFormation code. For this case, this is helpful because you have common infrastructure components that are shared across environments.


The Fn::ImportValue function is used to import values that have been exported in another stack. Since the networking team exports the VPC and subnet information, this can be used in the CloudFormation stack to reference those values.

upvoted 2 times

  **fanq10** 1 year ago

B is WRONG, you cannot use `TemplateURL` to retrieve Network Stack export values.

upvoted 4 times

  **sb333** 1 year, 1 month ago

B is incorrect. One of the reasons is that intrinsic functions are not allowed in the Parameters section. <https://repost.aws/knowledge-center/cloudformation-template-validation>

upvoted 2 times

  **bakamon** 1 year, 2 months ago

Selected Answer: C

C is the correct answer

upvoted 2 times

  **lunt** 1 year, 3 months ago

Selected Answer: C

C ipsingh is absolutely correct



upvoted 2 times

  **2pk** 1 year, 3 months ago

Selected Answer: C

Im 50/50 C or B, but B doesn't provide a clear approach for retrieving the exported values and placing position of Parameters section of the root template, which is not required to place it there, it must declare inside resource. So i think Answer C make sense.

upvoted 2 times

  **bcx** 1 year, 3 months ago

The template URL in B makes it wrong IMHO. You import the values from an exiting template importing the parameter exported by it.

upvoted 1 times

  **ParagSanyashiv** 1 year, 3 months ago

Selected Answer: C

C makes more sense


upvoted 2 times

  **meisme** 1 year, 4 months ago

Selected Answer: C

c is correct

upvoted 2 times

  **haazybanj** 1 year, 4 months ago

Selected Answer: B

B. Use nested stacks to define common infrastructure components. To access the exported values, use TemplateURL to reference the networking team's template. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root template. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.

This approach is a good fit because it allows the developer to define reusable infrastructure components as nested stacks. To retrieve VPC and subnet values, the intrinsic function `Fn::ImportValue` is used in the Parameters section of the root template, which retrieves the values from the output of the networking team's CloudFormation stack. To update existing environments, the `CreateChangeSet` and `ExecuteChangeSet` commands are used, which provides a way to easily update the deployed infrastructure. Additionally, the use of nested stacks helps to ensure consistency across environments.

upvoted 1 times

🗨️ 👤 **sb333** 1 year, 1 month ago

B is incorrect. One of the reasons is that intrinsic functions are not allowed in the Parameters section. <https://repost.aws/knowledge-center/cloudformation-template-validation>

upvoted 2 times

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

Selected Answer: B

Option B is correct. Using nested stacks, the common infrastructure components can be defined in separate templates that can be referenced by the root template. This allows for easy updates and maintenance of the common components. The networking team's CloudFormation template can be used to export the VPC and subnet values, which can be referenced in the root template using `Fn::ImportValue` intrinsic functions in the Parameters section. The `CreateChangeSet` and `ExecuteChangeSet` commands can be used to update the existing development environments.

Option C is not the best choice because using `Fn::ImportValue` intrinsic functions with the resources of the nested stack can lead to circular dependencies and make it difficult to manage the infrastructure.

upvoted 2 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

Selected Answer: C

C is correct

upvoted 2 times

A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.

Which solution will accomplish this?

- A. Create an AWS CloudFormation template that defines an AWS Inspector rule to check whether EBS encryption is enabled. Save the template to an Amazon S3 bucket that has been shared with all accounts within the company. Update the account creation script pointing to the CloudFormation template in Amazon S3.
- B. Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CLI. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization.
- C. Create an SCP in Organizations. Set the policy to prevent the launch of Amazon EC2 instances without encryption on the EBS volumes using a conditional expression. Apply the SCP to all AWS accounts. Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an `ec2:RunInstances` action.
- D. Deploy an IAM role to all accounts from a single trusted account. Build a pipeline with AWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the account. Publish a report to Amazon S3.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **YucelFuat** 6 months, 1 week ago

Selected Answer: B

Exam Tip -> Compliance = AWS Config
upvoted 1 times

🗳️ 👤 **dzn** 1 year ago

Selected Answer: B

Deploy CloudFormation template with encrypted-volumes in the `ConfigRuleName` property, AWS Config will automatically scan the environment and check for unencrypted EBS volumes.
upvoted 4 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

B is correct

upvoted 1 times

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

Selected Answer: B

B is the only solution meeting the criteria.
upvoted 3 times

🗳️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: B

B. Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CLI. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization, will accomplish the compliance check on all accounts.

Option A is incorrect because an AWS Inspector rule is used to analyze the behavior of the application on the EC2 instance, not to check the encryption of the EBS volume.
upvoted 3 times

🗳️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: B

B is right
upvoted 2 times

🗳️ 👤 **alce2020** 1 year, 11 months ago

Selected Answer: B

B is the answer
upvoted 2 times

A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations. Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts.

A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector.

Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

- A. Verify that AWS Systems Manager Agent is installed and is running on the EC2 instances that Amazon Inspector is not scanning.
- B. Associate the target EC2 instances with security groups that allow outbound communication on port 443 to the AWS Systems Manager service endpoint.
- C. Grant `inspector:StartAssessmentRun` permissions to the IAM role that the DevOps engineer is using.
- D. Configure EC2 Instance Connect for the EC2 instances that Amazon Inspector is not scanning.
- E. Associate the target EC2 instances with instance profiles that grant permissions to communicate with AWS Systems Manager.
- F. Create a managed-instance activation. Use the Activation Code and the Activation ID to register the EC2 instances.

Suggested Answer: ABC

Community vote distribution

ABE (100%)

  **Dimidrol**  1 year, 5 months ago

Selected Answer: ABE

A b e <https://docs.aws.amazon.com/inspector/latest/user/scanning-ec2.html>



upvoted 6 times

  **dzn**  6 months, 3 weeks ago

Selected Answer: ABE

C is not a fundamental solution. Because Inspector is actually able to run, and it is not the same IAM role that DevOps uses.

upvoted 3 times

  **thanhnv142** 7 months, 1 week ago

ABE are correct: Check if SSM agent is installed, check connection and permission of Ec2 that allows access to SSM

C: no need to grant `inspector:StartAssessmentRun` permissions because the dev has already finish the scanning task

D: There is not EC2 instance Connect, only need SSM agent

F: there is no managed-instance activation

upvoted 4 times

  **yorkicurke** 9 months ago

Selected Answer: ABE

the following link explains it all;

<https://repost.aws/knowledge-center/systems-manager-ec2-instance-not-appear>



upvoted 3 times

  **madperro** 1 year, 2 months ago

Selected Answer: ABE

ABE seem to be prerequisites to work with SSM and Inspector.

upvoted 2 times

  **bcx** 1 year, 3 months ago

Selected Answer: ABE

A B E is the correct one IMHO

upvoted 2 times

  **ParagSanyashiv** 1 year, 4 months ago

Selected Answer: ABE

ABE makes more sense.

upvoted 2 times

🗨️ 👤 **alice2020** 1 year, 4 months ago

A,B,E are correct <https://docs.aws.amazon.com/inspector/latest/user/scanning-ec2.html>

upvoted 3 times

🗨️ 👤 **jqso234** 1 year, 4 months ago

Selected Answer: ABE

Option C suggests granting `inspector:StartAssessmentRun` permissions to the IAM role being used by the DevOps engineer. However, this may not be relevant to the issue of instances not being scanned by Amazon Inspector, as the IAM role may already have the necessary permissions by default.

Therefore, A, B, E is a better choice in this case as it includes the necessary steps to ensure that the instances can communicate with AWS Systems Manager, which is required for Amazon Inspector to scan the instances.

upvoted 4 times

A development team uses AWS CodeCommit for version control for applications. The development team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy for CI/CD infrastructure. In CodeCommit, the development team recently merged pull requests that did not pass long-running tests in the code base. The development team needed to perform rollbacks to branches in the codebase, resulting in lost time and wasted effort.


A DevOps engineer must automate testing of pull requests in CodeCommit to ensure that reviewers more easily see the results of automated tests as part of the pull request review.

What should the DevOps engineer do to meet this requirement?

- A. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
- B. Create an Amazon EventBridge rule that reacts to the pullRequestCreated event. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.
- C. Create an Amazon EventBridge rule that reacts to pullRequestCreated and pullRequestSourceBranchUpdated events. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
- D. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.

Suggested Answer: D

Community vote distribution



MarDog Highly Voted 1 year, 8 months ago

C. Look at #3 in the below.

<https://container-devsecops.awssecworkshops.com/04-testing/>

upvoted 16 times

Gomer 9 months, 1 week ago

Link is dead

upvoted 1 times

madperro Highly Voted 1 year, 8 months ago

Selected Answer: B

B, we need to run tests only when pull request is created and we need to publish test results, not only badge.

upvoted 11 times

Igallard Most Recent 1 month, 3 weeks ago

Selected Answer: C

C covers both event scenarios: PR creation and update (when updating the source branch, for example, with a new commit).

upvoted 1 times

Gomer 9 months, 1 week ago

Selected Answer: C

"Automated Code Review on Pull Requests using AWS CodeCommit and AWS CodeBuild"

"The solution comprises of the following components:"

"Amazon EventBridge: AWS service to receive pullRequestCreated and pullRequestSourceBranchUpdated events and trigger Amazon EventBridge rule."

<https://aws.amazon.com/blogs/devops/automated-code-review-on-pull-requests-using-aws-codecommit-and-aws-codebuild/>

upvoted 3 times

seetpt 10 months, 2 weeks ago

Selected Answer: C

i go with C

upvoted 1 times

🗨️ 👤 **zijo** 11 months, 3 weeks ago

C is the answer to ensure code reviewers more easily see the results of automated tests as part of the pull request review
pullRequestStatusChanged event is triggered whenever the status of a pull request changes. This could include transitions like: Open to Closed (pull request is merged or marked as closed)
Closed to Open (pull request is reopened)
pullRequestCreated event is triggered whenever a new pull request is created in a CodeCommit repository.
pullRequestSourceBranchUpdated event is triggered whenever there are updates (new commits) pushed to the source branch of an open pull request

upvoted 4 times

🗨️ 👤 **shammous** 7 months, 2 weeks ago

I agree the C is the closed correct answer but it doesn't mention pullRequestStatusChanged (not sure why you mention it in your comment).

"The primary events in AWS CodeCommit that can trigger the pipeline are:

- pullRequestCreated: This event occurs when a new pull request is created.
- pullRequestSourceBranchUpdated: This event occurs when the source branch of an existing pull request is updated (e.g. when new commits are pushed to the branch)."

The other events that might be considered but I would exclude are:

- pullRequestStatusChanged: This event occurs when the status of a pull request changes, from closed to open (which we can consider), or from open to close (which we shouldn't consider in our case).
- pullRequestMerged: This event occurs when a pull request is merged. I would exclude it because we are looking to test before merging.

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

B is correct: we need to react when there is merge request (pullRequestCreated event)

A: we need to react when there is merge request, not when the status of merge request is changed (pullRequestStatusChanged event)

C: we only need to react when there is merge request, not when a sourcebranch is updated (pullRequestSourceBranchUpdated events)

D: we need to react when there is merge request, not when the status of merge request is changed (pullRequestStatusChanged event)

upvoted 2 times

🗨️ 👤 **gg_robin** 8 months, 1 week ago

If the source is updated after the PR is created, you don't run any tests against those changes.

upvoted 4 times

🗨️ 👤 **DucSiu** 1 year, 2 months ago

Why not B?

upvoted 2 times

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

Answer is C: the pullRequestStatusChanged only has two values (OPEN|CLOSED) so If there is any update in the code the tests will not run.

https://docs.aws.amazon.com/codecommit/latest/APIReference/API_PullRequestStatusChangedEventMetadata.html

upvoted 2 times

🗨️ 👤 **DZ_Ben** 1 year, 4 months ago

Selected Answer: C

I'll go for C. Tbh, I don't think we will need a lambda here as the event rule can definitely trigger the code pipeline & code build.

upvoted 2 times

🗨️ 👤 **RVivek** 1 year, 6 months ago

Selected Answer: C

<https://aws.amazon.com/blogs/devops/automated-code-review-on-pull-requests-using-aws-codecommit-and-aws-codebuild/>

upvoted 4 times

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

Correct C.

upvoted 2 times

🗨️ 👤 **vhernan** 1 year, 7 months ago

Selected Answer: C

C

run tests on pull requests created and when source branch receives new commits to re-run tests

upvoted 2 times

🗨️ 👤 **lunt** 1 year, 7 months ago

Selected Answer: C

Not sure why so much discussion. Triggers Rule: A CloudWatch Event Rule is triggered based on the following events: pullRequestSourceBranchUpdated or pullRequestCreated.

C is only viable option. I mean it even tells you the answer in the question "development team needed to perform rollbacks to branches in the codebase".

Ans is C.

upvoted 2 times

🗨️ 👤 **tartarus23** 1 year, 8 months ago

Selected Answer: C

This approach allows testing whenever a pull request is created or the source branch of a pull request is updated. When the tests are complete, the AWS Lambda function posts the test status badge as a comment on the pull request, providing visual feedback to reviewers directly in the context of the pull request review.

It's important to note that CodeBuild creates a build badge that provides status about the last build, which might not directly reflect the test results of the specific pull request. Posting the test results would provide more accurate and relevant information but doing so might require additional scripting or tooling not described in the available options.

upvoted 3 times

🗨️ 👤 **bcx** 1 year, 9 months ago

Selected Answer: C

C is the correct IMHO.

A pull request is just a branch that the requestor is asking to be merged in master/main. When you create a pull request you set the branch, that is the start, you have to use the current contents of the branch to execute the tests. When time passes developers add commits to that branch or force-push it, changing the contents of the PR's branch. That is the moment in which you have to trigger the tests. The PR comments and discussions may change, but that does not change the code so no need to perform new tests.

You only test when the PR is created and every time the branch is pushed (updated).

upvoted 3 times

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

Why not B?

upvoted 2 times

A company has deployed an application in a production VPC in a single AWS account. The application is popular and is experiencing heavy usage. The company's security team wants to add additional security, such as AWS WAF, to the application deployment. However, the application's product manager is concerned about cost and does not want to approve the change unless the security team can prove that additional security is necessary.

The security team believes that some of the application's demand might come from users that have IP addresses that are on a deny list. The security team provides the deny list to a DevOps engineer. If any of the IP addresses on the deny list access the application, the security team wants to receive automated notification in near real time so that the security team can document that the application needs additional security. The DevOps engineer creates a VPC flow log for the production VPC.

Which set of additional steps should the DevOps engineer take to meet these requirements MOST cost-effectively?

- A. Create a log group in Amazon CloudWatch Logs. Configure the VPC flow log to capture accepted traffic and to send the data to the log group. Create an Amazon CloudWatch metric filter for IP addresses on the deny list. Create a CloudWatch alarm with the metric filter as input. Set the period to 5 minutes and the datapoints to alarm to 1. Use an Amazon Simple Notification Service (Amazon SNS) topic to send alarm notices to the security team.
- B. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture all traffic and to send the data to the S3 bucket. Configure Amazon Athena to return all log files in the S3 bucket for IP addresses on the deny list. Configure Amazon QuickSight to accept data from Athena and to publish the data as a dashboard that the security team can access. Create a threshold alert of 1 for successful access. Configure the alert to automatically notify the security team as frequently as possible when the alert threshold is met.
- C. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture accepted traffic and to send the data to the S3 bucket. Configure an Amazon OpenSearch Service cluster and domain for the log files. Create an AWS Lambda function to retrieve the logs from the S3 bucket, format the logs, and load the logs into the OpenSearch Service cluster. Schedule the Lambda function to run every 5 minutes. Configure an alert and condition in OpenSearch Service to send alerts to the security team through an Amazon Simple Notification Service (Amazon SNS) topic when access from the IP addresses on the deny list is detected.
- D. Create a log group in Amazon CloudWatch Logs. Create an Amazon S3 bucket to hold query results. Configure the VPC flow log to capture all traffic and to send the data to the log group. Deploy an Amazon Athena CloudWatch connector in AWS Lambda. Connect the connector to the log group. Configure Athena to periodically query for all accepted traffic from the IP addresses on the deny list and to store the results in the S3 bucket. Configure an S3 event notification to automatically notify the security team through an Amazon Simple Notification Service (Amazon SNS) topic when new objects are added to the S3 bucket.

Suggested Answer: A

Community vote distribution

A (100%)

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

Selected Answer: A

A meets the requirements at the lowest cost.

upvoted 8 times


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

Selected Answer: A

opensearch cost a lot of \$\$\$\$. Athena got tons of things to do afterwards. It's used purely for interactive query.

natively, vpcflow sends logs to s3 or cloudwatch logs. no brainer answer

upvoted 5 times

 **thanhv142** Most Recent 7 months, 1 week ago

A is correct: push all VPC flow log to cloudwatch logs. Create metric filter to find denied IP addresses. Create cloudwatch alarm with the metric filter as input. Alarm's action is send noti to Security team via SNS

B: "Configure the alert to automatically notify the security team": alert cannot notify by itself. Must use SNS

C: This option uses both S3 bucket and "Amazon OpenSearch Service cluster" to store log files, which would cost a lot of money and unnecessary

D: This option uses both S3 bucket and VPC flow log to store log files, which is costly

upvoted 2 times

 **ele** 1 year, 3 months ago

Selected Answer: A

A most cost effective

upvoted 2 times

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

Selected Answer: A

To meet the requirements most cost-effectively, the DevOps engineer should create a log group in Amazon CloudWatch Logs and configure the VPC flow log to capture accepted traffic and to send the data to the log group. Then, create an Amazon CloudWatch metric filter for IP addresses on the deny list and create a CloudWatch alarm with the metric filter as input. Set the period to 5 minutes and the datapoints to alarm to 1. Finally, use an Amazon Simple Notification Service (Amazon SNS) topic to send alarm notices to the security team.

Option A is the correct answer. It provides a cost-effective solution that meets the requirements. The CloudWatch alarm notifies the security team in near real-time when traffic from an IP address on the deny list is detected. This will help the security team document that the application needs additional security. This solution only requires the use of AWS services that the company is already using, and does not require any additional services or tools.

upvoted 4 times

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

D. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.

upvoted 1 times

🗨️ 👤 **BaburTurk** 1 year ago

Option D does not mention an event bridge rule

upvoted 1 times

🗨️ 👤 **zoltthar_z** 9 months, 3 weeks ago

This answer is for a previous question

upvoted 1 times

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

The best way to automate testing of pull requests in CodeCommit is to use Amazon EventBridge rules to detect pullRequestStatusChanged events, which are triggered when a pull request's status changes. When this event occurs, the DevOps engineer can create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Finally, program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete. This approach ensures that reviewers more easily see the results of automated tests as part of the pull request review, without the need to perform manual testing or rollbacks in the codebase.

upvoted 2 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

Selected Answer: A

A seems correct

upvoted 2 times

A DevOps engineer has automated a web service deployment by using AWS CodePipeline with the following steps:

1. An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
2. An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
3. A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment.

The quality assurance (QA) team requests permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal penetration testing tool to conduct manual tests. The tool will be invoked by a REST API call. Which combination of actions should the DevOps engineer take to fulfill this request? (Choose two.)

- A. Insert a manual approval action between the test actions and deployment actions of the pipeline.
- B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
- C. Update the CodeDeploy deployment groups so that they require manual approval to proceed.
- D. Update the pipeline to directly call the REST API for the penetration testing tool.
- E. Update the pipeline to invoke an AWS Lambda function that calls the REST API for the penetration testing tool.

Suggested Answer: AE

Community vote distribution

AE (78%)

AD (22%)

🗳️ **tartarus23** Highly Voted 1 year, 8 months ago

Selected Answer: AE

Explanation:

The manual approval action (A) will allow the QA team to inspect the build artifact and run their internal penetration testing tool before the deployment to the production environment proceeds.

Using an AWS Lambda function (E) would provide an automated way to call the REST API of the penetration testing tool. This would allow for the tests to be conducted automatically within the pipeline. This is beneficial because it ensures consistency in the testing process and could be run programmatically, reducing manual steps.

upvoted 8 times

🗳️ **iulian0585** Most Recent 7 months, 1 week ago

Selected Answer: AE

Option D (updating the pipeline to directly call the REST API for the penetration testing tool) is not recommended because it tightly couples the pipeline with the QA team's tool, making it less flexible and harder to maintain. Using a Lambda function as an intermediary provides better separation of concerns and easier maintainability.

upvoted 2 times

🗳️ **jamesf** 7 months, 2 weeks ago

Selected Answer: AE

Should be AE

Although there are limitation 15mins of Lambda function.

But Option D is wrong as CodePipeline does not have the ability to execute HTTP requests "directly".

<https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-invoke-lambda-function.html>

upvoted 1 times

🗳️ **jamesf** 7 months, 1 week ago

For option A, keywords: conduct manual tests

upvoted 1 times

🗳️ **zijo** 11 months, 3 weeks ago

This is tricky but AD should be a better choice because of the 15 min timeout of Lambda functions. To call REST API in CodePipeline these are the two options

For complex API calls, security requirements, and access to external resources, an AWS Lambda function is the recommended approach.

For simple API calls with limited requirements, consider the inline script approach within CodeBuild, but with caution due to security and maintainability limitations.

upvoted 1 times

🗨️ 👤 **Shasha1** 1 year ago

AE there is no way to call REST API directly in the code pipeline, it is possible invoke via Lambda function only

upvoted 4 times

🗨️ 👤 **dzn** 1 year ago

Selected Answer: AE

CodePipeline does not have the ability to execute HTTP requests "directly".

upvoted 3 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

A and D are correct: a manual approval action between the test actions and deployment actions allows tester to verify and test built artifacts before allowing deploying to production

B and C: no mentions of test and deployment env

E: manual test take more than 15 minutes, which is the maximum execution time of lambda

upvoted 1 times

🗨️ 👤 **a54b16f** 1 year, 1 month ago

Selected Answer: AE

D is wrong, alternative option (not using Lambda, for example, if the pen testing will take more than 15 minutes) is using codebuild, either add a new codebuild for pen testing, or update existing unit testing codebuild to include pen testing. You should never run Pen testing inside codepipeline directly , it lacks the hooks to collect test result, inform result, etc

upvoted 2 times

🗨️ 👤 **2pk** 1 year, 4 months ago

Selected Answer: AD

Tricky one:

CodeDeploy can't do actions directly like invoke REST API but code Build can.

e.g. it's mentioned to test build artifacts.. So after the build artifact is created This means the solution uses Code Build even not from Code Build you can setup a python script and run it directly using Code Build Command:

I'd not use Lambda as an alternative due to the time taken for penetration tests would take more than 15 mins. and the pipeline would failed with Lambda execution timeout.

upvoted 4 times

🗨️ 👤 **shammous** 7 months, 2 weeks ago

The lambda function would just invoke the REST API, it won't execute the pen test itself. An asynchronous mechanism involving SQS could handle the waiting time between the requesting sending and the response receiving, which can indeed last more than 15mn.

upvoted 1 times

🗨️ 👤 **Seoyong** 1 year, 6 months ago

conducting manual tests might takes more than 15m.

upvoted 1 times

🗨️ 👤 **s50600822** 1 year, 7 months ago

A,

E in practice is a cheap and handy off-switch, recommended, for some contributors to CI/CD that we don't control directly. However, no idea what the writer of the question wanted.

upvoted 1 times

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

why don't you choose D

upvoted 1 times

🗨️ 👤 **habros** 1 year, 8 months ago

Selected Answer: AE

I'll choose AE. I can tie up multiple REST calls in a Lambda and customize it as I wish. A web hook is not flexible in this aspect I feel.

upvoted 2 times

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

Selected Answer: AD



AD, Lambda is not needed, a webhook can call REST API directly.

upvoted 1 times

🗨️ 👤 **cocegas** 1 year, 6 months ago

But there is no option to invoke call an API directly = <https://docs.aws.amazon.com/codepipeline/latest/userguide/integrations-action-type.html#integrations-invoke>

upvoted 1 times



  **bcx** 1 year, 9 months ago

Selected Answer: AE

"AWS Lambda is a compute service that lets you run code without provisioning or managing servers. You can create Lambda functions and add them as actions in your pipelines. Because Lambda allows you to write functions to perform almost any task, you can customize the way your pipeline works. "

<https://docs.aws.amazon.com/codepipeline/latest/userguide/actions-invoke-lambda-function.html>



upvoted 3 times

  **qsergii** 1 year, 9 months ago

Selected Answer: AD

A & D, lambda (E) is extra and not needed.

upvoted 2 times

  **Akaza** 1 year, 9 months ago

Selected Answer: AE

Yepp A, E for me

upvoted 2 times

A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to scale up to handle all traffic. How should a DevOps engineer meet these requirements?


- A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session data. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
- B. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session data. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
- C. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS for PostgreSQL with cross-region replication for session data. Deploy the web application with client-side logic to call the API Gateway directly.
- D. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session data. Enable an Amazon CloudFront weighted distribution across regions. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Suggested Answer: A

Community vote distribution

A (81%)

D (19%)

 **davdan99** Highly Voted 1 year, 2 months ago

Selected Answer: A

I think it is A, We can have failover with CloudFront, but it can't have weighted routing, Here is the link of how automatic failover works in the CloudFront

<https://disaster-recovery.workshop.aws/en/services/networking/cloudfront/cloudfront-origin-group.html>

upvoted 7 times

 **jamesf** Most Recent 7 months, 2 weeks ago


A

Keywords: web application - ElasticBeanstalk, weighted routing required.

DynamoDB Global Table required.

As understand, Cloudfront not support weighted routing.

upvoted 3 times

 **auxwww** 7 months, 2 weeks ago

Selected Answer: A

A - correct - Elasticbeanstalk - option of ALB to register route53 with active-active alias with health checks and weighted routing

"In active-active failover, all the records that have the same name, the same type (such as A or AAAA), and the same routing policy (such as weighted or latency) are active unless Route 53 considers them unhealthy. Route 53 can respond to a DNS query using any healthy record."

D - incorrect because no ALB in front of ASG.

upvoted 1 times

 **Gomer** 9 months, 1 week ago

Selected Answer: D

If the requirement is for "1% of requests should route to the secondary region to continuously", then that means the secondary region is in continuously in an Active state (Active/Active). A "request" is not a health check. You also have to have auto-scaling to dynamically pick up any extra traffic. The question is a little weird, in I don't know you dynamically adjust the weighted routing policy to steer all traffic to the secondary region. I just know that "D" is the closest choice to meeting the specified requirements. This is absolutely an "Active-Active" design using weighted routing at some level, and auto-scaling just meets the demand wherever it comes from. I think "latency-based" routing would make more sense, but the requirements are clearly describing "weighted routing" and Active-Active design.

<https://aws.amazon.com/blogs/networking-and-content-delivery/latency-based-routing-leveraging-amazon-cloudfront-for-a-multi-region-active-active-architecture/>

upvoted 1 times

🗨️ **Gomer** 9 months, 1 week ago

Just to clarify, the scenarios is absolutely desirvinb "weighted routing" with 99% to 1% traffic split between regions for normal operation (unbalanced Active/Active).

upvoted 1 times

🗨️ **seetpt** 10 months, 2 weeks ago

Selected Answer: A

A is ok

upvoted 1 times

🗨️ **DanShone** 11 months, 4 weeks ago

Selected Answer: A

A is correct

B + C no DynamoDB Global Tables

D - does not use Route53

upvoted 2 times

🗨️ **thanhnv142** 1 year, 1 month ago

A is correct: beanstalk is literally designed for this specific purpose

upvoted 2 times

🗨️ **Jaguaroooo** 1 year, 2 months ago

It is A, 1% of the traffic should be going to the 2ndary site. so that's weighted routing.

upvoted 2 times

🗨️ **DucSiu** 1 year, 2 months ago

Why not B?

upvoted 1 times

🗨️ **davdan99** 1 year, 2 months ago

We have to use DynamoDB Global tables for make db acessable from 2 regions.

upvoted 1 times

🗨️ **dencccc** 1 year, 4 months ago

It's A

upvoted 2 times

🗨️ **rahulsingha2112** 1 year, 4 months ago

A is correct answer

upvoted 1 times

🗨️ **ekki** 1 year, 4 months ago

Answer is D.

"Testing Regional failover"

<https://aws.amazon.com/blogs/networking-and-content-delivery/latency-based-routing-leveraging-amazon-cloudfront-for-a-multi-region-active-active-architecture/>

upvoted 1 times

🗨️ **zijo** 11 months, 3 weeks ago

The title of this page mentions Active-Active scenario and not Active-Passive as mentioned in this question.

upvoted 1 times

🗨️ **BaburTurk** 1 year, 6 months ago

Selected Answer: D

Option A uses Elastic Beanstalk, which is not as scalable as Auto Scaling groups.

D is correct

upvoted 1 times

🗨️ **BaburTurk** 1 year, 4 months ago

A- Route 53 does offer the capability to automatically route traffic to the secondary region in case of a disruption. In the context of the requirements specified, option A seems to be a feasible solution as it involves the use of AWS Elastic Beanstalk for deployment, DynamoDB global tables for session data replication, and a weighted routing policy in Route 53 for traffic distribution across regions. The health checks can ensure that traffic is routed to the secondary region automatically in case of a disruption in the main region.

Therefore, considering the ability of Route 53 to automatically reroute traffic, option A appears to be the most appropriate solution for meeting the specified disaster recovery requirements.

upvoted 2 times

🗨️ 👤 **mamila** 1 year, 7 months ago

The answer is NONE of them, none of them specified both weighted and failover routing policies.

upvoted 1 times

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

Selected Answer: A

D is not offering scaling on DRP. A offers scaling by using BeanStalk.

upvoted 2 times

🗨️ 👤 **csG13** 1 year, 8 months ago

Selected Answer: A

Going for A given that DynamoDB global tables can replicate data across selected regions in near real-time. Clearly weighting and failover, thus Route53 should be selected.

It's not D because Cloudfront cannot do weighted routing.

upvoted 4 times

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

Selected Answer: D

A and D are very similar but with using different services (BeanStalk vs CloudFront). However A is using R53 traffic distribution and D is using CF traffic distribution. I think D is better in this case. Note that not all applications will run easily on BeanStalk too.

upvoted 3 times

A company runs an application on Amazon EC2 instances. The company uses a series of AWS CloudFormation stacks to define the application resources. A developer performs updates by building and testing the application on a laptop and then uploading the build output and CloudFormation stack templates to Amazon S3. The developer's peers review the changes before the developer performs the CloudFormation stack update and installs a new version of the application onto the EC2 instances.

The deployment process is prone to errors and is time-consuming when the developer updates each EC2 instance with the new application. The company wants to automate as much of the application deployment process as possible while retaining a final manual approval step before the modification of the application or resources.

The company already has moved the source code for the application and the CloudFormation templates to AWS CodeCommit. The company also has created an AWS CodeBuild project to build and test the application.

Which combination of steps will meet the company's requirements? (Choose two.)



- A. Create an application group and a deployment group in AWS CodeDeploy. Install the CodeDeploy agent on the EC2 instances.
- B. Create an application revision and a deployment group in AWS CodeDeploy. Create an environment in CodeDeploy. Register the EC2 instances to the CodeDeploy environment.
- C. Use AWS CodePipeline to invoke the CodeBuild job, run the CloudFormation update, and pause for a manual approval step. After approval, start the AWS CodeDeploy deployment.
- D. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step. After approval, run the CloudFormation change sets and start the AWS CodeDeploy deployment.
- E. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step. After approval, start the AWS CodeDeploy deployment.

Suggested Answer: BD

Community vote distribution

AD (68%)



BD (32%)

  **tartarus23** Highly Voted 1 year, 8 months ago

Selected Answer: AD



(A) This step sets up the environment to use AWS CodeDeploy for application deployments. CodeDeploy uses an agent installed on the EC2 instances to perform the deployment tasks.

(D) This option uses CodePipeline to orchestrate the process. CodeBuild is used to build and test the application. CloudFormation is used to prepare the infrastructure updates as change sets. A manual approval step is inserted before applying the changes. After approval, the CloudFormation change sets are applied, and then CodeDeploy is invoked to deploy the new version of the application to the EC2 instances.
upvoted 15 times


  **ky11223344** Highly Voted 1 year, 6 months ago

Selected Answer: BD

There is no application group in CodeDeploy
upvoted 5 times

  **fanq10** 1 year, 6 months ago

- EC2 needs to install the CodeDeploy agent.
- CodeDeploy does not need to register EC2 instances, instead of it uses tag filter.
Therefore, B is incorrect, A is correct. Final answer: AD
upvoted 8 times

  **Karamen** 1 year, 4 months ago

@fanq10

you are right.

CodeDeploy doesn't require registering EC2 instances, it filters by tag
upvoted 3 times

  **jamesf** Most Recent 7 months, 2 weeks ago

Selected Answer: AD

AD

A - To run CodeDeploy on EC2, need agent.

D - The approval step will trigger both CloudFormation and CodeDeploy

B incorrect as no mention of installing agent on EC2 and CodeDeploy doesn't require registering EC2 instances, it filters by tag.

C incorrect as The approval step does not affect CloudFormation updates, which is not accepted

E incorrect as The approval step only allows CodeDeploy but no have CloudFormation updates

upvoted 2 times

  **alex_heavy** 8 months ago

Selected Answer: BD

B:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/codedeploy-agent.html>

You can configure automatic installation and updates of the CodeDeploy agent when you create your deployment group in the console.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/applications.html>



After you configure instances, but before you can deploy a revision, you must create an application in CodeDeploy. An application is simply a name or container used by CodeDeploy to ensure the correct revision, deployment configuration, and deployment group are referenced during a deployment.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/application-revisions.html>

In CodeDeploy, a revision contains a version of the source files CodeDeploy will deploy to your instances or scripts CodeDeploy will run on your instances.

<https://aws.amazon.com/ru/blogs/devops/using-codedeploy-environment-variables/>

upvoted 2 times


  **xdkonorek2** 8 months, 2 weeks ago

Selected Answer: AD

for those who vote B: what is creating environment in code deploy?

I think application group means application and registering instances with code deploy is basically creating deployment group, and instance is not registered manually it has to be tagged

upvoted 1 times

  **seetpt** 10 months, 2 weeks ago

Selected Answer: AD

AD is ok

upvoted 1 times

  **4555894** 1 year ago

Selected Answer: AD

A- <https://docs.aws.amazon.com/codedeploy/latest/userguide/codedeploy-agent.html>

D - This option correctly utilizes AWS CodePipeline to invoke the CodeBuild job and create CloudFormationchange sets. It adds a manual approval step before executing the change sets and starting the AWSCodeDeploy deployment. This ensures that the deployment process is automated while retaining the final manual approval step.

upvoted 3 times

  **Shasha1** 1 year ago

AD

Needs to install code deploy agent and give necessary permission for access S3 bucket where it will be stored the application revision. then EC2 instance will download the application revision from the S3 bucket. Therefore Answer A is correct. if we use System manager only the EC2 instance can be installed and updated automatically.

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

A and D are correct: A - To run codedeploy on EC2, need agent. D - The approval step will trigger both cloudformation and codedeploy

B: no mention of installing agent on EC2

C: The approval step doesnot affect CloudFormation updates, which is not accepted

E: The approval step only allows codedeploy but not CloudFormation updates

upvoted 2 times

  **Jaguaroooo** 1 year, 2 months ago

AD is the correct answer, you need the CD agent in order to use code deploy. And D is correct also because you can do your testing during codebuild and finally do a change set review and then approval.

upvoted 1 times

🗲️ 👤 **2pk** 1 year, 4 months ago

Selected Answer: BD

Code deploy agent installation can be skipped when you setting up Code Deploy group .. e.g.

You can configure automatic installation and updates of the CodeDeploy agent when you create your deployment group in the console.

Why A is wrong - cause there is no application group only deployment group and when setting up deployment group you can setup agent installation automatically.

[https://docs.aws.amazon.com/codedeploy/latest/userguide/instances-ec2-](https://docs.aws.amazon.com/codedeploy/latest/userguide/instances-ec2-create.html#:~:text=Note-,You%20can%20configure%20automatic%20installation%20and%20updates%20of%20the%20CodeDeploy%20agent%20when%20Did%20this%20page)

[create.html#:~:text=Note-,You%20can%20configure%20automatic%20installation%20and%20updates%20of%20the%20CodeDeploy%20agent%20when%20Did%20this%20page](https://docs.aws.amazon.com/codedeploy/latest/userguide/instances-ec2-create.html#:~:text=Note-,You%20can%20configure%20automatic%20installation%20and%20updates%20of%20the%20CodeDeploy%20agent%20when%20Did%20this%20page)

upvoted 4 times

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

i agree with you on the explanation that there are no application group. but how can you use code deploy to push code to the EC2's if they have no agent installed?

upvoted 1 times

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

Selected Answer: AD

AD is right.

upvoted 2 times

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

The CodeDeploy agent must be installed on your Amazon EC2 instance before using it in CodeDeploy deployments

upvoted 1 times

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

There is nothing like application group in code deploy answer is BD

upvoted 2 times

🗲️ 👤 **ele** 1 year, 10 months ago

Selected Answer: AD

AD need codedeploy agent, and review CFN changes set, before run update

upvoted 2 times

🗲️ 👤 **PhuocT** 1 year, 10 months ago

Selected Answer: AD

A and D are correct

upvoted 2 times

🗲️ 👤 **alce2020** 1 year, 10 months ago

Selected Answer: BD

B and D are correct

upvoted 1 times

A DevOps engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The engineer needs to implement a deployment strategy that:


- Launches a second fleet of instances with the same capacity as the original fleet.
- Maintains the original fleet unchanged while the second fleet is launched.
- Transitions traffic to the second fleet when the second fleet is fully deployed.
- Terminates the original fleet automatically 1 hour after transition.

Which solution will satisfy these requirements?

- A. Use an AWS CloudFormation template with a retention policy for the ALB set to 1 hour. Update the Amazon Route 53 record to reflect the new ALB.
- B. Use two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new one. Create an application version lifecycle policy to terminate the original environment in 1 hour.
- C. Use AWS CodeDeploy with a deployment group configured with a blue/green deployment configuration. Select the option Terminate the original instances in the deployment group with a waiting period of 1 hour.
- D. Use AWS Elastic Beanstalk with the configuration set to Immutable. Create an .ebextension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

Suggested Answer: C

Community vote distribution



haazybanj Highly Voted 1 year, 10 months ago

Option B, using two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new one, would not launch a second fleet of instances. Instead, it would create a new environment and deploy the application version to it. It also requires the use of an application version lifecycle policy to terminate the original environment in 1 hour.

Option D, using AWS Elastic Beanstalk with the configuration set to Immutable and creating an .ebextension to set the deletion policy of the ALB to 1 hour, would not launch a second fleet of instances, and it would not maintain the original fleet unchanged while the second fleet is launched. Additionally, the .ebextension approach is not the recommended way to delete resources in AWS.

Therefore, the correct option is C, using AWS CodeDeploy with a deployment group configured with a blue/green deployment configuration and selecting the option to Terminate the original instances in the deployment group with a waiting period of 1 hour.

upvoted 12 times

hayjaykay Most Recent 1 month ago

Selected Answer: C

Key words: EC2 for CodeDeploy, Elastic Beanstalk for CodePipeline.

upvoted 2 times

newpotato 5 months, 1 week ago

Option B would require more manual intervention and configuration, making it a less optimal solution compared to the seamless blue/green deployment and automatic fleet termination provided by Option C.

Option D involves unnecessary complexity around setting ALB deletion policies, and while immutable deployments offer zero-downtime updates, they don't fully meet the core requirements of automatic traffic shifting and fleet termination

upvoted 1 times

73d8cc9 9 months, 1 week ago

Selected Answer: D

Immutable strategy with Elastic Beanstalk involves deploying additional instance while Blue/Green strategy involves deploying another environment. The key difference is environment vs instances

upvoted 1 times

seetpt 10 months, 2 weeks ago

Selected Answer: C

C is ok

upvoted 1 times

🗨️ 👤 **vmahilevskyi** 1 year ago

Selected Answer: C

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-groups-create-blue-green.html>

upvoted 3 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

A is correct: The question ask for a solution to automatic deployment of EC2 instances, which is the job of cloudFormation

- B and D is irrelevant because it is use to deploy webapps only, not EC2 instances

- C is also irrelevant because codedeploy (literlly by the name: CODEdeploy) is only used for deploying code, not EC2 instances, which is not code. Dont know why ChatGPT recommend this, but it is wrong definitely

upvoted 1 times

🗨️ 👤 **jojom19980** 1 year ago

No C is the more accurate and logical

upvoted 1 times

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

Selected Answer: C

C looks like the best solution.

upvoted 3 times

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

Selected Answer: C

To satisfy the requirements of launching a second fleet of instances with the same capacity as the original fleet, maintaining the original fleet unchanged while the second fleet is launched, transitioning traffic to the second fleet when the second fleet is fully deployed, and terminating the original fleet automatically 1 hour after the transition, the best solution is to use AWS CodeDeploy with a blue/green deployment configuration, and selecting the option to Terminate the original instances in the deployment group with a waiting period of 1 hour.

upvoted 3 times

🗨️ 👤 **alce2020** 1 year, 11 months ago

Selected Answer: C

Option C

upvoted 2 times

A video-sharing company stores its videos in Amazon S3. The company has observed a sudden increase in video access requests, but the company does not know which videos are most popular. The company needs to identify the general access pattern for the video files. This pattern includes the number of users who access a certain file on a given day, as well as the number of pull requests for certain files. How can the company meet these requirements with the LEAST amount of effort?

- A. Activate S3 server access logging. Import the access logs into an Amazon Aurora database. Use an Aurora SQL query to analyze the access patterns.
- B. Activate S3 server access logging. Use Amazon Athena to create an external table with the log files. Use Athena to create a SQL query to analyze the access patterns.
- C. Invoke an AWS Lambda function for every S3 object access event. Configure the Lambda function to write the file access information, such as user, S3 bucket, and file key, to an Amazon Aurora database. Use an Aurora SQL query to analyze the access patterns.
- D. Record an Amazon CloudWatch Logs log message for every S3 object access event. Configure a CloudWatch Logs log stream to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL application. Perform a sliding window analysis.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **zijo** 5 months, 2 weeks ago

Bis the Answer because Athena is designed for these type of use cases

AWS Athena is a serverless interactive query service that lets you analyze data stored in Amazon Simple Storage Service (S3) using standard SQL.

upvoted 4 times

🗳️ 👤 **thanhnv142** 7 months, 1 week ago

B is correct: Use S3 in combination with Athena is the recommended way to analyze data

A: setups of Aurora is complex and unnecessary. It also more costly than B

C and D are both too complicated.

upvoted 2 times

🗳️ 👤 **habros** 1 year, 2 months ago

Selected Answer: B

B is so much simpler. Athena can do interactive queries on S3 data.

upvoted 4 times

🗳️ 👤 **madperro** 1 year, 2 months ago

Selected Answer: B

B is correct and simplest.

upvoted 2 times

🗳️ 👤 **tom_uk** 1 year, 3 months ago

Selected Answer: B

B is the answer

upvoted 1 times

🗳️ 👤 **bcx** 1 year, 3 months ago

Selected Answer: B

B is the natural way to do it

upvoted 1 times

🗳️ 👤 **gdtypk** 1 year, 3 months ago

Selected Answer: B

<https://repost.aws/ja/knowledge-center/analyze-logs-athena>

upvoted 1 times

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

Selected Answer: B

B is right

upvoted 1 times

🗨️ 👤 **Sazeka** 1 year, 4 months ago

I would go with B

upvoted 1 times

🗨️ 👤 **alce2020** 1 year, 4 months ago

Selected Answer: B

B is correct


upvoted 1 times

A development team wants to use AWS CloudFormation stacks to deploy an application. However, the developer IAM role does not have the required permissions to provision the resources that are specified in the AWS CloudFormation template. A DevOps engineer needs to implement a solution that allows the developers to deploy the stacks. The solution must follow the principle of least privilege. Which solution will meet these requirements?

- A. Create an IAM policy that allows the developers to provision the required resources. Attach the policy to the developer IAM role.
- B. Create an IAM policy that allows full access to AWS CloudFormation. Attach the policy to the developer IAM role.
- C. Create an AWS CloudFormation service role that has the required permissions. Grant the developer IAM role a `cloudformation:*` action. Use the new service role during stack deployments.
- D. Create an AWS CloudFormation service role that has the required permissions. Grant the developer IAM role the `iam:PassRole` permission. Use the new service role during stack deployments.

Suggested Answer: B

Community vote distribution



fuzzycom 8 months, 2 weeks ago

D is totally correct
upvoted 1 times

4555894 1 year ago

Selected Answer: D

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html>
upvoted 2 times

thanhnv142 1 year, 1 month ago

D is correct: Need to create a role for Cloud formation that has the required permissions. Then adding `iam:PassRole` permission to the dev IAM role to allow them to pass this role to CF

A: no mention of creating the required permissions for ACF. Additionally, should not grant permissions for dev.

B: grant full access is against the least privilege policy

C: no mention of granting `iam:PassRole` permission to the dev

upvoted 4 times

imymoco 1 year, 2 months ago

A is incorrect; A would also allow resources to be used from outside of cfn.

Therefore, D is correct.

upvoted 1 times

jason7 1 year, 6 months ago

Selected Answer: D

Option D allows you to create a dedicated AWS CloudFormation service role with the precise permissions required for stack deployments. Then, you grant the developer IAM role the `iam:PassRole` permission, which enables it to pass the service role to AWS CloudFormation without granting it broad IAM permissions. This approach aligns best with the principle of least privilege and ensures developers can deploy stacks while maintaining control over their permissions.

upvoted 2 times

ogwu2000 1 year, 7 months ago

B is the answer. DC wrong - Nothing like CloudFormation service-role.

upvoted 1 times

fanq10 1 year, 6 months ago

B is not best practice of using CloudFormation.

D is correct, 100% sure. `iam:PassRole` to a CloudFormation Service Role (take a look at this:`

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html#:~:text=you%20can%20use.,Important,-When%20you%20specify>

upvoted 1 times

🗨️ **DZ_Ben** 1 year, 4 months ago

Should be D! See here <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html>

upvoted 1 times

🗨️ **madperro** 1 year, 8 months ago

Selected Answer: D

D is the right answer.

upvoted 1 times

🗨️ **tartarus23** 1 year, 8 months ago

Selected Answer: D

This solution follows the principle of least privilege by creating a specific AWS CloudFormation service role that only has the permissions required for the resources in the AWS CloudFormation stack. The developers are then granted permission to pass this role (iam:PassRole) to the AWS CloudFormation service when they initiate stack deployments, which allows the service to act on behalf of the developer to provision the specified resources.

upvoted 1 times

🗨️ **bcx** 1 year, 9 months ago

Selected Answer: D

D, you pass the role that can create the resources, the user does not have the right to create the resources himself but can pass the role to CloudFormation so CloudFormation assumes it. IMHO.

upvoted 1 times

🗨️ **2pk** 1 year, 10 months ago

Selected Answer: D

This allows them to provision the required resources specified in the CloudFormation template without granting them full access to AWS CloudFormation or the underlying resources.

upvoted 1 times

🗨️ **ele** 1 year, 10 months ago

Selected Answer: D

D, passrole is right action

upvoted 1 times

🗨️ **gdtypk** 1 year, 10 months ago

Selected Answer: D

https://docs.aws.amazon.com/ja_jp/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

upvoted 3 times

🗨️ **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: D

D is more suitable in this case.

upvoted 1 times

🗨️ **Frodo_the_cat** 1 year, 10 months ago

C. Create an AWS CloudFormation service role that has the required permissions. Grant the developer IAM role a cloudformation:* action. Use the new service role during stack deployments.

By creating an AWS CloudFormation service role with the required permissions, the DevOps engineer can control the resources that the developers can access. This approach ensures that the developers have only the necessary permissions to deploy the stacks, without granting them excessive permissions that could be exploited by malicious actors. The IAM policy granting a cloudformation:* action to the developer IAM role allows the developers to use the AWS CloudFormation service role and deploy the stacks with the required resources.

Option A, creating an IAM policy that allows the developers to provision the required resources, is not a good solution because it could potentially grant the developers too much access to resources they don't need. This violates the principle of least privilege.

upvoted 1 times

🗨️ **Frodo_the_cat** 1 year, 10 months ago

Option B, creating an IAM policy that allows full access to AWS CloudFormation, is not a good solution either, as it grants excessive permissions to the developers.

Option D, creating an AWS CloudFormation service role with the required permissions and granting the developer IAM role the iam:PassRole permission, allows the developers to assume the service role and deploy the stacks with the required resources. However, this option grants additional permissions to the developer IAM role, which could be abused by malicious actors

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

upvoted 1 times

🗨️ 👤 **kassem77** 1 year, 10 months ago

D it is

upvoted 1 times

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

Selected Answer: D

Option D is the recommended solution to meet the requirements because it follows the principle of least privilege. The IAM policy that allows the developers to provision the required resources should be created and associated with the IAM role, which should be assigned the iam:PassRole permission for the AWS CloudFormation service role. By doing so, the IAM role can only assume the specific AWS CloudFormation service role and deploy the stack with the required permissions, and not have full access to all resources or full access to AWS CloudFormation.

upvoted 1 times

🗨️ 👤 **alce2020** 1 year, 10 months ago

Selected Answer: B

B it is

upvoted 2 times

A production account has a requirement that any Amazon EC2 instance that has been logged in to manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with the Amazon CloudWatch Logs agent configured. How can this process be automated?

- A. Create a CloudWatch Logs subscription to an AWS Step Functions application. Configure an AWS Lambda function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a second Lambda function once a day that will terminate all instances with this tag.
- B. Create an Amazon CloudWatch alarm that will be invoked by the login event. Send the notification to an Amazon Simple Notification Service (Amazon SNS) topic that the operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
- C. Create an Amazon CloudWatch alarm that will be invoked by the login event. Configure the alarm to send to an Amazon Simple Queue Service (Amazon SQS) queue. Use a group of worker instances to process messages from the queue, which then schedules an Amazon EventBridge rule to be invoked.
- D. Create a CloudWatch Logs subscription to an AWS Lambda function. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a daily Lambda function that terminates all instances with this tag.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Aesthet** 7 months ago

Option D: "with this tag"

So, there will be one tag, like ShoudTerminate: true. But by doing so we will terminate ALL instances with a tag - even those created 10 minutes ago. It doesn't seem correct, or am I missing something?

upvoted 1 times

🗳️ 👤 **fuzzycom** 8 months, 2 weeks ago

D is best answer.

hint: question includes "~~Amazon CloudWatch Logs agent configured"

Lambda function is keyword.

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

D is correct:

A: If using step function, no need to include "Amazon EventBridge rule to invoke a second Lambda function"

B: With this method, policy-breaching Ec2 would be terminated manually, which cannot ensure that they are terminated within 24 hours

C: no mention of terminating the instances

upvoted 4 times

🗳️ 👤 **imymoco** 1 year, 2 months ago

D is correct; with B, SNS can cause delays.

upvoted 2 times

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

Selected Answer: D

D is the best answer.

upvoted 1 times

🗳️ 👤 **haazybanj** 1 year, 10 months ago

Selected Answer: D

D. Create a CloudWatch Logs subscription to an AWS Lambda function. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a daily Lambda function that terminates all instances with this tag.

upvoted 3 times

🗳️ 👤 **alce2020** 1 year, 10 months ago

Selected Answer: D

D is the correct answer

upvoted 3 times

A company has enabled all features for its organization in AWS Organizations. The organization contains 10 AWS accounts. The company has turned on AWS CloudTrail in all the accounts. The company expects the number of AWS accounts in the organization to increase to 500 during the next year. The company plans to use multiple OUs for these accounts.

The company has enabled AWS Config in each existing AWS account in the organization. A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization.

Which solution will meet this requirement?

- A. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call. Configure the rule to invoke an AWS Lambda function that enables trusted access to AWS Config for the organization.
- B. In the organization's management account, create an AWS CloudFormation stack set to enable AWS Config. Configure the stack set to deploy automatically when an account is created through Organizations.
- C. In the organization's management account, create an SCP that allows the appropriate AWS Config API calls to enable AWS Config. Apply the SCP to the root-level OU.
- D. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call. Configure the rule to invoke an AWS Systems Manager Automation runbook to enable AWS Config for the account.

Suggested Answer: B

Community vote distribution

B (100%)

 **thanhnv142** 7 months, 1 week ago

B is correct: The question ask a solution to "enables AWS Config automatically" for all future accounts. In AWS org, to provision or configure resources on other accounts, we use ACF

A, C and D: no mention of ACF

upvoted 2 times

 **thanhnv142** 7 months, 1 week ago

A: trusted access to AWS Config: this is used by other services to access to AWS config, not for account

D: enable AWS Config for the account: this means we only activate AWS config for the management account, not the newly created ones

upvoted 1 times


 **2pk** 10 months, 1 week ago

B:

Details the new feature with enable trusted access to new accounts in any region

<https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-cloudformation.html>

upvoted 1 times

 **hje0329** 1 year ago

B

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

upvoted 1 times

 **madperro** 1 year, 2 months ago

Selected Answer: B

B is the best solution.

upvoted 1 times

 **samgyeopsal** 1 year, 3 months ago

B

<https://aws.amazon.com/about-aws/whats-new/2020/02/aws-cloudformation-stacksets-introduces-automatic-deployments-across-accounts-and-regions-through-aws-organizations/>

upvoted 2 times



 **haazybanj** 1 year, 4 months ago

Selected Answer: B

The correct solution to enable AWS Config automatically for all future AWS accounts created in the organization is Option B: In the organization's management account, create an AWS CloudFormation stack set to enable AWS Config. Configure the stack set to deploy automatically when an account is created through Organizations.

Option C is incorrect because although it suggests creating an SCP that allows the appropriate AWS Config API calls to enable AWS Config and applying the SCP to the root-level OU, it does not specifically enable AWS Config automatically for all future AWS accounts that are created in the organization.

upvoted 4 times

  **Olelukoe** 8 months, 3 weeks ago

In terms of Option C: SCP can only Deny access, not Allow

upvoted 3 times

  **alce2020** 1 year, 4 months ago

Selected Answer: B

B is correct


upvoted 1 times

A company has many applications. Different teams in the company developed the applications by using multiple languages and frameworks. The applications run on premises and on different servers with different operating systems. Each team has its own release protocol and process. The company wants to reduce the complexity of the release and maintenance of these applications. The company is migrating its technology stacks, including these applications, to AWS. The company wants centralized control of source code, a consistent and automatic delivery pipeline, and as few maintenance tasks as possible on the underlying infrastructure. What should a DevOps engineer do to meet these requirements?

- A. Create one AWS CodeCommit repository for all applications. Put each application's code in a different branch. Merge the branches, and use AWS CodeBuild to build the applications. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- B. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- C. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time and to create one AMI for each server. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 fleets by using these AMIs.
- D. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build one Docker image for each application in Amazon Elastic Container Registry (Amazon ECR). Use AWS CodeDeploy to deploy the applications to Amazon Elastic Container Service (Amazon ECS) on infrastructure that AWS Fargate manages.

Suggested Answer: D

Community vote distribution



haazybanj Highly Voted 1 year, 4 months ago

Selected Answer: D

Option D is the best choice to meet the requirements of centralized control of source code, a consistent and automatic delivery pipeline, and minimal maintenance tasks.

Option D is the best choice because it allows each application to have its own repository and build process, but uses containerization to create a consistent and automatic delivery pipeline that can be easily deployed to Amazon ECS on infrastructure that AWS Fargate manages. This approach also provides scalability and ease of maintenance.

upvoted 11 times

thanhnv142 Most Recent 7 months, 1 week ago

D is correct: "centralized control of source code" = CodeCommit. "Consistent and automatic delivery pipeline" = codepipeline/codebuilde/codedeploy. "as few maintenance tasks as possible on the underlying infrastrucutr" = containerization

A: "CodeCommit repository for all applications": should not, need separate repos for each app

B and C: no mention of containerization (fargate, ECS)

upvoted 1 times

DaddyDee 11 months, 1 week ago

D is the Answer: <https://aws.amazon.com/blogs/compute/building-deploying-and-operating-containerized-applications-with-aws-fargate/>

upvoted 1 times

ddedqdw 1 year, 1 month ago

Selected Answer: A

A ISNT CORRECT?

upvoted 1 times

davdan99 8 months ago

Of course no, it is the most wrong one, It saying to have all the applications code in one repo (bad thing to do), separated in branches, and after that merge them (second very bad thing to do).

upvoted 2 times

RickSk 1 year, 2 months ago

Option D.

I was torn between C and D, but the requirement for ease of maintenance on the underlying infrastructure clearly points to ECS.

upvoted 3 times

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

Selected Answer: D

D is the best option, there is virtually no infrastructure to manage.

upvoted 1 times

🗨️ 👤 **ele** 1 year, 3 months ago

Selected Answer: D

D is best option.

upvoted 1 times

A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions.


Which combination of actions should be taken to address the latency issues? (Choose three.)

- A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
- B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
- C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
- D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
- E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
- F. Convert the DynamoDB table to a global table.

Suggested Answer: CDF

Community vote distribution

CDF (100%)


 **haazybanj** Highly Voted 1 year, 4 months ago

Selected Answer: CDF

C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.

D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the nearest healthy ALB, based on the latency between the user and the ALBs.

F. Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application
upvoted 10 times

 **xdkonorek2** Most Recent 4 months, 3 weeks ago

Technically converting dynamodb table to global table requires creating replica in another region with cross-region replication and you don't "convert" you add a replica in "global tables" in specified region so this answers are a little bit misleading.

Probably F is better than A since they name this operation as "converting" e.g. here <https://aws.amazon.com/blogs/aws/new-convert-your-single-region-amazon-dynamodb-tables-to-global-tables/>
upvoted 2 times

 **thanhnv142** 7 months, 1 week ago

D is, of course, correct: <apply a core set of security controls to an existing set of AWS accounts> and <The accounts are in an organization in AWS Organizations> means we need ACF template to deploy these set of security controls. <Individual account administrators must not be able to edit or delete any of the baseline resources> means we need scp to deny permission

A and B: no mention of SCP

C: this option deploy the rules by AWS Config management account, which is not correct because we need ACF. Additionally, no mention of denying modification to CloudTrail trails
upvoted 1 times

 **thanhnv142** 7 months, 1 week ago

CDF: <opened a new office on a different continent> and <The users in the new office are experiencing high latency> means they need to replicate their existing site to the new region. <Amazon EC2 instances behind an Application Load Balancer (ALB)> means they need to replicate both these. <address the latency issues> means they need route53 with latency-based and health check

A: <cross-Region replication> is used for backup only, not a live site. It would introduce a lot of latency

B: <Auto Scaling group global resources>: there is no such thing

E: No mention of latency-based.

upvoted 1 times

🗨️ 👤 **z_inderjot** 8 months, 2 weeks ago

Selected Answer: CDF

CDF very easy

upvoted 1 times

🗨️ 👤 **|||||** 10 months, 3 weeks ago

CDF is collect

upvoted 1 times

🗨️ 👤 **alice2020** 1 year, 4 months ago

Selected Answer: CDF

C,D,F are correct

upvoted 2 times


A DevOps engineer needs to apply a core set of security controls to an existing set of AWS accounts. The accounts are in an organization in AWS Organizations. Individual teams will administer individual accounts by using the AdministratorAccess AWS managed policy. For all accounts, AWS CloudTrail and AWS Config must be turned on in all available AWS Regions. Individual account administrators must not be able to edit or delete any of the baseline resources. However, individual account administrators must be able to edit or delete their own CloudTrail trails and AWS Config rules.

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

- A. Create an AWS CloudFormation template that defines the standard account resources. Deploy the template to all accounts from the organization's management account by using CloudFormation StackSets. Set the stack policy to deny Update/Delete actions.
- B. Enable AWS Control Tower. Enroll the existing accounts in AWS Control Tower. Grant the individual account administrators access to CloudTrail and AWS Config.
- C. Designate an AWS Config management account. Create AWS Config recorders in all accounts by using AWS CloudFormation StackSets. Deploy AWS Config rules to the organization by using the AWS Config management account. Create a CloudTrail organization trail in the organization's management account. Deny modification or deletion of the AWS Config recorders by using an SCP.
- D. Create an AWS CloudFormation template that defines the standard account resources. Deploy the template to all accounts from the organization's management account by using CloudFormation StackSets. Create an SCP that prevents updates or deletions to CloudTrail resources or AWS Config resources unless the principal is an administrator of the organization's management account.

Suggested Answer: C

Community vote distribution



haazybanj Highly Voted 1 year, 10 months ago

Selected Answer: C

C

This solution meets the requirements in the most operationally efficient way. It uses AWS CloudFormation StackSets to deploy AWS Config recorders in all accounts and AWS Config rules to the organization, which can be centrally managed from an AWS Config management account. A CloudTrail organization trail can also be created in the organization's management account to collect logs from all accounts. An SCP can be used to deny modification or deletion of the AWS Config recorders, ensuring that the baseline resources cannot be modified or deleted by individual account administrators. However, individual account administrators can still edit or delete their own CloudTrail trails and AWS Config rules.

upvoted 19 times

koenigParas2324 1 year, 3 months ago

this solution lacks clarity on allowing individual account administrators control over their CloudTrail trails.

upvoted 4 times

bnagaraja9099 1 year, 2 months ago

C is good.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

upvoted 1 times

bnagaraja9099 1 year, 2 months ago

An SCP restricts permissions for IAM users and roles in member accounts, including the member account's root user. Any account has only those permissions permitted by every parent above it. If a permission is blocked at any level above the account, either implicitly (by not being included in an Allow policy statement) or explicitly (by being included in a Deny policy statement), a user or role in the affected account can't use that permission, even if the account administrator attaches the AdministratorAccess IAM policy with */* permissions to the user.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

upvoted 1 times

a1234321606 1 year, 3 months ago

Why C? If you deny modification or deletion of the AWS Config recorders by using an SCP, how do individual account administrators edit or delete their own CloudTrail trails and AWS Config rules?

upvoted 5 times

dzn Highly Voted 1 year ago

Selected Answer: B

When Control Tower is enabled, AWS-GR_CLOUDTRAIL_ENABLED and AWS-GR_CONFIG_ENABLED will enable CloudTrail and Config in all available regions. The guardrails are automatically set to disallow changes to baseline resources.

A, C, D - No mention about baseline resource.

upvoted 6 times

  **lovekiller** **Most Recent** 1 month ago

Selected Answer: D

For the people fighting over Option C or D:

Option D is a better choice because it deploys all the required baseline resources—both CloudTrail and AWS Config—across every account using CloudFormation StackSets, and then locks them down with an SCP. This setup makes sure that only the management account's administrators can change or remove these core security controls.

Meanwhile, individual account admins can still manage any extra CloudTrail trails or Config rules they create on their own.

But, Option C, mainly focuses on AWS Config recorders and doesn't clearly cover CloudTrail, so it does not offer us as complete a solution.

upvoted 2 times

  **Slays** 2 months, 1 week ago

Selected Answer: D



Option D:

Create an AWS CloudFormation Template: Develop a template that defines the standard resources, including CloudTrail and AWS Config, configured to operate in all available AWS Regions.

Deploy Using CloudFormation StackSets: Utilize AWS CloudFormation StackSets from the organization's management account to deploy the template across all member accounts. This approach ensures consistent configuration and simplifies management.

Implement a Service Control Policy (SCP): Establish an SCP that restricts updates or deletions of CloudTrail and AWS Config resources. This policy should allow only the organization's management account administrators to perform such actions, preventing individual account administrators from making unauthorized changes.

upvoted 2 times

  **youonebe** 2 months, 3 weeks ago

Selected Answer: D

This is the most operationally efficient solution. Using CloudFormation StackSets ensures standard resources are consistently deployed, and SCPs provide the necessary restrictions and flexibility.

upvoted 2 times

  **steli0** 3 months, 2 weeks ago

Selected Answer: C

I think C because the SCP defines the principal being an administrator from the management account, not the individual account.

upvoted 1 times

  **steli0** 3 months, 2 weeks ago

moreover Principals are not supported in SCPs

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html#scp-syntax-unsupported

upvoted 1 times

  **BrusingWayne** 3 months, 2 weeks ago

D. CloudFormation StackSets + SCP with conditional permissions:

Centralized deployment of resources

SCP prevents modifications to core resources

Allows admins to edit their own resources (by implication)

Matches all requirements efficiently

upvoted 3 times

  **rk0509** 6 months, 4 weeks ago

Selected Answer: D

D is correct

upvoted 4 times

🗨️ **jamesf** 7 months, 2 weeks ago

Selected Answer: C

I think should be C

Keywords: "an existing set of AWS accounts"

upvoted 1 times

🗨️ **trungtd** 7 months, 4 weeks ago

Selected Answer: D

must be D

upvoted 2 times

🗨️ **seetpt** 10 months, 1 week ago

Selected Answer: C

I agree with C

upvoted 1 times

🗨️ **Mordans** 11 months, 3 weeks ago

Selected Answer: C

Option C is the most operationally efficient and meets all the requirements: ensuring CloudTrail and AWS Config are enabled in all regions, preventing the deletion or editing of baseline resources by individual account administrators, while still allowing them the flexibility to manage their own specific resources. This approach uses centralized control mechanisms (AWS Config management account and organization trail for CloudTrail) and leverages SCPs for enforcement, aligning with best practices for security and governance in AWS Organizations.

upvoted 2 times

🗨️ **CloudHandsOn** 1 year ago

Selected Answer: D

Im going with D. SCPs is what helps us here

upvoted 2 times

🗨️ **vn_thanhtung** 10 months, 1 week ago

but SCP not support direct principal.

upvoted 1 times

🗨️ **vmahilevskyi** 1 year ago

Selected Answer: D

D for me.

I think C is incorrect because "However, individual account administrators must be able to edit or delete their own CloudTrail trails and AWS Config rules." requirement is not satisfied because this answer has nothing about individual account administrators are able to edit their own CloudTrail trails. Organisational trail can be edited only from management or delegated administrator account.

upvoted 5 times

🗨️ **[Removed]** 1 year ago

Selected Answer: C

C for sure

upvoted 2 times

🗨️ **thanhnv142** 1 year ago

Selected Answer: D

D is correct: This denies modifications to AWS config or cloudtrail unless the principal is the management account

A: No explicitly mention of denying modifications to Config or cloudtrail

B: No explicitly mention of denying modifications to Config or cloudtrail

C: < Create a CloudTrail organization trail in the organization's management account>: This means the deny rule only affects the management account

upvoted 2 times

🗨️ **Chelseajcole** 1 year, 1 month ago

Selected Answer: D

C is using AWS Config Recorder, AWS Config uses the configuration recorder to detect changes in your resource configurations and capture these changes as configuration items.

It is not used for prevent you doing something, it is detecting something

upvoted 3 times

A company has its AWS accounts in an organization in AWS Organizations. AWS Config is manually configured in each AWS account. The company needs to implement a solution to centrally configure AWS Config for all accounts in the organization. The solution also must record resource changes to a central account.

Which combination of actions should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Configure a delegated administrator account for AWS Config. Enable trusted access for AWS Config in the organization.
- B. Configure a delegated administrator account for AWS Config. Create a service-linked role for AWS Config in the organization's management account.
- C. Create an AWS CloudFormation template to create an AWS Config aggregator. Configure a CloudFormation stack set to deploy the template to all accounts in the organization.
- D. Create an AWS Config organization aggregator in the organization's management account. Configure data collection from all AWS accounts in the organization and from all AWS Regions.
- E. Create an AWS Config organization aggregator in the delegated administrator account. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

Suggested Answer: BD

Community vote distribution

AE (84%)

BD (16%)

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

Selected Answer: AE

AE

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/>

A - When enabling trust - the service-linked role will be created but not the other way around.

E - the delegated account will be the account that manages AWS config so it should collect all data centrally.

upvoted 17 times

jamesf **Most Recent** 7 months, 2 weeks ago

Selected Answer: AE

A - You can enable trusted access using either the AWS Config console or the AWS Organizations console.

<https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

upvoted 1 times

zijo 11 months, 3 weeks ago

AE is the answer

AWS Config offers an organization-wide data aggregation capability called the Config organization aggregator. It allows you to collect and view configuration data from all member accounts within your AWS Organization in a single location. This centralizes your view of resource configurations and compliance posture across your entire AWS environment.

upvoted 1 times

thanhnv142 1 year, 1 month ago

A and E are correct: <AWS Config is manually configured in each AWS account> means we don't need ACF (only used for the deployment of AWS config). <centrally configure AWS Config for all accounts> means we need to allow a central account to control AWS config in all member accounts.

- <record resource changes to a central account> means we need to collect data from all member accounts and push to the central account

B: service-linked role only used for interacting with other AWS services

C: no need ACF

D: we need AWS Config organization aggregator in the delegated administrator account, not the organization's management account

upvoted 1 times

hoaille257 1 year, 5 months ago

Selected Answer: AE

AE is most correct

upvoted 2 times

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

Selected Answer: AE

Here you have the Tutorial :)

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/>

upvoted 3 times

🗄️ 👤 **rhinozD** 1 year, 9 months ago

Selected Answer: AE

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/>

<https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

upvoted 3 times

🗄️ 👤 **Kodoma** 1 year, 9 months ago

BE is the most efficient

upvoted 3 times

🗄️ 👤 **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: BD

BD is most suitable in this case

upvoted 3 times

🗄️ 👤 **2pk** 1 year, 4 months ago

Why ? it says setup service linked role in management account not in Delegated account?

upvoted 1 times

🗄️ 👤 **jqso234** 1 year, 11 months ago

Selected Answer: BD

The correct answers are B and D. Option B is correct because it suggests configuring a delegated administrator account for AWS Config and creating a service-linked role for AWS Config in the organization's management account. This allows AWS Config to perform supported operations within the accounts in the organization, and enables trusted access. Option D is correct because it suggests creating an AWS Config organization aggregator in the organization's management account and configuring data collection from all AWS accounts in the organization and from all AWS Regions, which enables multi-account, multi-region data aggregation. Options A and E are not correct because they do not suggest using a service-linked role for AWS Config or creating an AWS Config organization aggregator in the organization's management account.

upvoted 2 times

🗄️ 👤 **Dimidrol** 1 year, 11 months ago

Selected Answer: AE

AE . <https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

upvoted 3 times

A company wants to migrate its content sharing web application hosted on Amazon EC2 to a serverless architecture. The company currently deploys changes to its application by creating a new Auto Scaling group of EC2 instances and a new Elastic Load Balancer, and then shifting the traffic away using an Amazon Route 53 weighted routing policy.

For its new serverless application, the company is planning to use Amazon API Gateway and AWS Lambda. The company will need to update its deployment processes to work with the new application. It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base.

Which deployment strategy will meet these requirements?

- A. Use AWS CDK to deploy API Gateway and Lambda functions. When code needs to be changed, update the AWS CloudFormation stack and deploy the new version of the APIs and Lambda functions. Use a Route 53 failover routing policy for the canary release strategy.
- B. Use AWS CloudFormation to deploy API Gateway and Lambda functions using Lambda function versions. When code needs to be changed, update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strategy. Promote the new version when testing is complete.
- C. Use AWS Elastic Beanstalk to deploy API Gateway and Lambda functions. When code needs to be changed, deploy a new version of the API and Lambda functions. Shift traffic gradually using an Elastic Beanstalk blue/green deployment.
- D. Use AWS OpsWorks to deploy API Gateway in the service layer and Lambda functions in a custom layer. When code needs to be changed, use OpsWorks to perform a blue/green deployment and shift traffic gradually.

Suggested Answer: B

Community vote distribution

B (100%)

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

Selected Answer: B

The deployment strategy that will meet the company's requirements is B. Use AWS CloudFormation to deploy API Gateway and Lambda functions using Lambda function versions. When code needs to be changed, update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strategy. Promote the new version when testing is complete.

Explanation:

Option B provides a deployment strategy for the company's new serverless architecture, allowing the company to retain the ability to test new features on a small number of users before rolling the features out to the entire user base. Using AWS CloudFormation, the company can deploy API Gateway and Lambda functions using Lambda function versions. When code needs to be changed, the company can update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strategy. Once testing is complete, the new version can be promoted.

upvoted 12 times

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


Selected Answer: B

A Wrong: not using Canary, or Blue/green

C Wrong: Beanstalk is not serverless deployment platform

D Wrong: Irrelevant, OpsWork is configuration management platform and situation is requesting application deployment /AWS resource provisioning platform

upvoted 7 times

 **beanxyz** 1 year, 6 months ago

Your answer is correct but the explanation is not.

A is wrong because we can't use Route 53 failover routing for canary release. If it says Route53 weighted routing, then it is a possible option.

D is wrong because when you use blue/green mode, the switch from blue to green is all done at once, not like a granular canary change

upvoted 3 times

 **teo2157** Most Recent 3 months, 1 week ago

Selected Answer: B

It could be either A or B but the key here is that "Use a Route 53 failover routing policy for the canary release strategy." is a wrong statement, it should be a Route 53 weighted routing policy so B is the correct answer

upvoted 1 times


  **thanhnv142** 1 year, 1 month ago

B is correct: <serverless architecture> means ECS, lambda, Beanstalk. < It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base> means canary deployment

A: <Route 53 failover routing policy for the canary release strategy>: there is no such thing

C and D: no mention of canary deployment

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

B is correct: <serverless architecture> means ECS, lambda, Beanstalk. < It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base> means canary deployment

A: <Route 53 failover routing policy for the canary release strategy>: there is no such thing

C: no mention of canary deployment

upvoted 1 times

  **Jeanphi72** 1 year, 10 months ago

Selected Answer: B

ONLY B is possible

upvoted 1 times

A development team uses AWS CodeCommit, AWS CodePipeline, and AWS CodeBuild to develop and deploy an application. Changes to the code are submitted by pull requests. The development team reviews and merges the pull requests, and then the pipeline builds and tests the application.


Over time, the number of pull requests has increased. The pipeline is frequently blocked because of failing tests. To prevent this blockage, the development team wants to run the unit and integration tests on each pull request before it is merged.

Which solution will meet these requirements?

- A. Create a CodeBuild project to run the unit and integration tests. Create a CodeCommit approval rule template. Configure the template to require the successful invocation of the CodeBuild project. Attach the approval rule to the project's CodeCommit repository.
- B. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit. Create a CodeBuild project to run the unit and integration tests. Configure the CodeBuild project as a target of the EventBridge rule that includes a custom event payload with the CodeCommit repository and branch information from the event.
- C. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit. Modify the existing CodePipeline pipeline to not run the deploy steps if the build is started from a pull request. Configure the EventBridge rule to run the pipeline with a custom payload that contains the CodeCommit repository and branch information from the event.
- D. Create a CodeBuild project to run the unit and integration tests. Create a CodeCommit notification rule that matches when a pull request is created or updated. Configure the notification rule to invoke the CodeBuild project.

Suggested Answer: B

Community vote distribution



haazybanj Highly Voted 1 year, 4 months ago

Selected Answer: B

To run the unit and integration tests on each pull request before it is merged, a solution that listens to pullRequestCreated events and runs a CodeBuild project to execute tests would be the most appropriate option.

Option B describes a solution that creates an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit and configures a CodeBuild project to run the unit and integration tests, passing the CodeCommit repository and branch information from the event as a custom payload.

Therefore, option B is the correct answer.

upvoted 12 times

that1guy Most Recent 3 months, 4 weeks ago

Selected Answer: C

These days it would be C instead of B, it's very common to reuse the same pipeline but with conditions to skip certain steps depending on the branch.

<https://aws.amazon.com/blogs/devops/aws-codepipeline-adds-support-for-branch-based-development-and-monorepos/>

upvoted 1 times

jojom19980 6 months, 3 weeks ago

Selected Answer: A

The Answer should Be A because this option is allows test the code and the approval is depending on test's result

upvoted 2 times

vn_thanhtung 3 months, 3 weeks ago

The development team reviews and merges the pull requests, and then the pipeline builds and tests the application.

upvoted 1 times

Jon_aws 1 month, 1 week ago

shouldn't review be after build and test?

upvoted 1 times

thanhnv142 7 months, 1 week ago

A is definitely correct: <The development team reviews and merges the pull requests> and <the development team wants to run the unit and integration tests on each pull request before it is merged> means the dev team always review all pull requests and they need a solution to test committed code before merging to main. option A allow them to do tests and manually approve it before allow merging

B C and D: no mention of the step that allow the dev team to manually approve the merge.

upvoted 2 times

🗨️ 👤 **2pk** 10 months ago

B is the answer . D is wrong. Code commit only can setup notification rule to SNS topics or Chatbot.

upvoted 2 times

🗨️ 👤 **BaburTurk** 1 year ago

Selected Answer: B

<https://aws.amazon.com/blogs/devops/validating-aws-codecommit-pull-requests-with-aws-codebuild-and-aws-lambda/>

upvoted 3 times

🗨️ 👤 **Seoyong** 1 year ago

Selected Answer: D

Only D covers create pull request and update pull request

upvoted 1 times

🗨️ 👤 **zoltar_z** 9 months, 3 weeks ago

Can't be D because you can trigger a Codebuild project with CodeCommit notification rules

upvoted 2 times

🗨️ 👤 **Aja1** 1 year, 1 month ago

Option B is the most appropriate solution as it uses Amazon EventBridge rules to automatically trigger a CodeBuild project for running tests on each pull request, enabling early testing and preventing pipeline blockages due to failing tests after merging.

upvoted 1 times

🗨️ 👤 **totopopo** 1 year, 1 month ago

Selected Answer: D

Nothing in requirements says that the team wants a blockage of the pull request merge.

And the B solution talks only about the "PullRequestCreated" which is not enough, it has to be reexecuted at any event on the pull request.

upvoted 3 times

🗨️ 👤 **Snape** 1 year, 1 month ago

Selected Answer: B

A wrong: Dev team will need to manually approve each pull request before merging, this can be time consuming and error-prone.

C Wrong: Modifying the existing codepipeline is not necessary

D wrong: No prevention from pipeline being blocked by failing tests

upvoted 3 times

A company has an application that runs on a fleet of Amazon EC2 instances. The application requires frequent restarts. The application logs contain error messages when a restart is required. The application logs are published to a log group in Amazon CloudWatch Logs. An Amazon CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service (Amazon SNS) topic when the logs contain a large number of restart-related error messages. The application engineer manually restarts the application on the instances after the application engineer receives a notification from the SNS topic.


A DevOps engineer needs to implement a solution to automate the application restart on the instances without restarting the instances.

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

- A. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instances. Configure the SNS topic to invoke the runbook.
- B. Create an AWS Lambda function that restarts the application on the instances. Configure the Lambda function as an event destination of the SNS topic.
- C. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instances. Create an AWS Lambda function to invoke the runbook. Configure the Lambda function as an event destination of the SNS topic.
- D. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instances. Configure an Amazon EventBridge rule that reacts when the CloudWatch alarm enters ALARM state. Specify the runbook as a target of the rule.

Suggested Answer: B

Community vote distribution



daburahjail Highly Voted 1 year, 5 months ago

Selected Answer: D

It is debatable, as both C and D are correct and simple in their own ways, however, take a look at the number of components in each approach:

C: CW -> SNS -> LAMBDA -> SSM (4)

D: CW -> EVENTBRIDGE -> SSM (3)

There is an extra component (SNS) to maintain on C, also, there is some coding involved on this option, which also needs to be maintained. Even if we already have the SNS created on option C, we still have to go there to remove the notification and configure the lambda invocation.

Option D has fewer components, and require less customization.

upvoted 22 times

ParagSanyashiv Highly Voted 1 year, 10 months ago

Selected Answer: C

C makes more sense here

upvoted 10 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: D

D is more simpler solution than C.

upvoted 1 times

xdkonorek2 8 months, 1 week ago

Selected Answer: D

D)

B is wrong since it's way easier to use SSM automation runbook to execute logic inside instance using "run command" action within automation runbook than doing this with lambda

upvoted 1 times

zijo 11 months, 3 weeks ago

A is not possible - AWS Systems Manager (SSM) Run Command or Automation runbooks cannot be directly triggered by an Amazon SNS topic. Then C and D are the next best options. C is flexible but D is the most simple solution

upvoted 1 times

🗨️ 👤 **Diego1414** 1 year ago

Selected Answer: D

Both C and D are valid answers. However, D is less complicated.

upvoted 3 times

🗨️ 👤 **jojom19980** 1 year ago

Selected Answer: D

C is correct , But D is more easy to implement , cost saving, managed services by AWS ^_^

upvoted 2 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

B is correct: <implement a solution to automate the application restart on the instances> means we need to automate the restart step. We can use lambda, AWS system manager. <CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service> means we already have the alarm. We just need to simply trigger the restart process with lambda

A, C and D are all too complicated compared to B. They ask for "the MOST operationally efficient manner", not the most complicated one

upvoted 1 times

🗨️ 👤 **vn_thanhtung** 10 months ago

Option B not correct.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/running-automations-event-bridge.html>

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

B is correct: <implement a solution to automate the application restart on the instances> means we need to automate the restart step. We can use lambda, AWS system manager. <CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service> means we already have the alarm. We just need to simply trigger the restart process with lambda

A, C and D are all too complicated compared to A. They ask for "the MOST operationally efficient manner", not the most complicated one

upvoted 1 times

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

Selected Answer: D

For me D is the answer , because we use lambda for the custom operations , if we already have SSM automation to perform that same action then why writing our custom logic in lambda ?

upvoted 2 times

🗨️ 👤 **csG13** 1 year, 2 months ago

Selected Answer: D

It's D. Here is a reference:

<https://aws.amazon.com/blogs/mt/use-amazon-eventbridge-rules-to-run-aws-systems-manager-automation-in-response-to-cloudwatch-alarms/>

upvoted 4 times

🗨️ 👤 **HugoFM** 1 year, 3 months ago

Selected Answer: D

D It's the most simples approach. But C its also a solution, but why build and mantain a lambda?

upvoted 3 times

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

Selected Answer: B

I think is B, you only need to create the lambda and update the SNS to the lambda,

upvoted 2 times

🗨️ 👤 **n1w** 1 year, 4 months ago

Selected Answer: B

B seems like the shortest number of steps given that SNS already exists

upvoted 2 times

🗨️ 👤 **AWSdeveloper08** 1 year, 5 months ago

Selected Answer: D

Ill go with D too, less components, less configurations

upvoted 3 times

🗨️ 👤 **beanxyz** 1 year, 6 months ago

Selected Answer: D

B is wrong because SSM document is used to run on managed instances so definitely more efficient than lambda.

C is wrong because although this solution should work, we need to write a lambda script to invoke the runbook, while in D we don't need to do it
upvoted 2 times

  **beanxyz** 1 year, 6 months ago

I think both C and D will work, but the question is to chose the most efficient way, so I pickup D.
upvoted 2 times

A DevOps engineer at a company is supporting an AWS environment in which all users use AWS IAM Identity Center (AWS Single Sign-On). The company wants to immediately disable credentials of any new IAM user and wants the security team to receive a notification. Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

- A. Create an Amazon EventBridge rule that reacts to an IAM CreateUser API call in AWS CloudTrail.
- B. Create an Amazon EventBridge rule that reacts to an IAM GetLoginProfile API call in AWS CloudTrail.
- C. Create an AWS Lambda function that is a target of the EventBridge rule. Configure the Lambda function to disable any access keys and delete the login profiles that are associated with the IAM user.
- D. Create an AWS Lambda function that is a target of the EventBridge rule. Configure the Lambda function to delete the login profiles that are associated with the IAM user.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic that is a target of the EventBridge rule. Subscribe the security team's group email address to the topic.
- F. Create an Amazon Simple Queue Service (Amazon SQS) queue that is a target of the Lambda function. Subscribe the security team's group email address to the queue.

Suggested Answer: ACE

Community vote distribution

ACE (100%)

🗳️ **mrjaehong** 3 months, 4 weeks ago

The IAM user that was created cannot have an access key from the beginning. You need to log in and get an access key.
upvoted 1 times

🗳️ **0b005fc** 10 months, 3 weeks ago

Took the test 4/15 and passed. Almost all of the questions appeared.
ACE is correct.
upvoted 1 times

🗳️ **thanhnv142** 1 year, 1 month ago

ACE are correct: <disable credentials of any new IAM user> means disable all access key and profile related to the user. <the security team to receive a notification> means SNS
B: GetLoginProfile API is not equal to creating new user
D: we should delete all access key and profile related to the user, not just profile
F: we need SNS, not SQS
upvoted 4 times

🗳️ **khchan123** 1 year, 1 month ago

Selected Answer: ACE

Answer ACE
upvoted 2 times

🗳️ **yuliaqwerty** 1 year, 2 months ago

Answer ACE
upvoted 1 times

🗳️ **Snape** 1 year, 8 months ago

Selected Answer: ACE

No Brainer
upvoted 3 times

🗳️ **Jeanphi72** 1 year, 10 months ago

Selected Answer: ACE

My answer ACE
upvoted 4 times

🗳️ **haazybanj** 1 year, 10 months ago

Selected Answer: ACE

ACE is the right answer
upvoted 3 times


A company wants to set up a continuous delivery pipeline. The company stores application code in a private GitHub repository. The company needs to deploy the application components to Amazon Elastic Container Service (Amazon ECS). Amazon EC2, and AWS Lambda. The pipeline must support manual approval actions.

Which solution will meet these requirements?

- A. Use AWS CodePipeline with Amazon ECS. Amazon EC2, and Lambda as deploy providers.
- B. Use AWS CodePipeline with AWS CodeDeploy as the deploy provider.
- C. Use AWS CodePipeline with AWS Elastic Beanstalk as the deploy provider.
- D. Use AWS CodeDeploy with GitHub integration to deploy the application.

Suggested Answer: B

Community vote distribution



haazybanj Highly Voted 1 year, 10 months ago

Selected Answer: B

B is correct

upvoted 7 times

Just_Ninja Highly Voted 1 year, 8 months ago

Selected Answer: B

Because the Term "The pipeline must support manual approval actions."

That is not possible without a pipeline :)

upvoted 7 times

1rob Most Recent 2 months, 1 week ago

Selected Answer: B

Lambda is not defined as a deployment provider. Only as an "invoke" option. Amazon ECS is possible as deploy provider , check <https://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html> where it gives: CodeDeploy is a deployment service that automates application deployments to Amazon EC2 instances, on-premises instances, serverless Lambda functions, or Amazon ECS services. . So I go for B.

upvoted 1 times

Zdujgfr567783ff 2 months, 2 weeks ago

Selected Answer: A

Option B is partially correct but lacks native support for ECS deployments. AWS CodeDeploy is excellent for deploying to EC2 and Lambda, but it doesn't natively handle ECS deployments without additional configuration.

upvoted 1 times

Zdujgfr567783ff 2 months, 2 weeks ago

Selected Answer: A

asked chat GPT says a

upvoted 1 times

hzaki 6 months, 1 week ago

A is correct, CodeDeploy can't deploy the ECS

upvoted 1 times

thanhnv142 1 year, 1 month ago

A is correct

upvoted 1 times

thanhnv142 1 year ago

Correction: B is correct

upvoted 1 times

due 1 year, 3 months ago

Selected Answer: B

The solution for deploy ECS by codePipeline and codeDeploy

Create your CodeDeploy application and deployment group (ECS compute platform

<https://docs.aws.amazon.com/codepipeline/latest/userguide/tutorials-ecs-ecr-codedeploy.html#tutorials-ecs-ecr-codedeploy-deployment>
upvoted 4 times

🗨️ 👤 **RVivek** 1 year, 5 months ago

Selected Answer: A

Why not A ?

B (Code depoly providr does not support ECS)

D does not have "codepipeleine" and the question says ""The pipeline must support manual approval actions."

So A is the only feasible option

upvoted 3 times

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

using codedeploy we can deploy to ecs and even can perform blue / green deployment. Codedeploy support all there of the deployment strategies

upvoted 3 times

🗨️ 👤 **Radeeka** 1 year, 6 months ago

Selected Answer: D

Answer D.

Source: <https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-partners-github.html>

The question is asking for a application code stored in a GitHub repository.

upvoted 1 times

🗨️ 👤 **Radeeka** 1 year, 6 months ago

My bad, Above only support EC2 and OnPrem.

upvoted 1 times

🗨️ 👤 **Aja1** 1 year, 7 months ago

option A with AWS CodePipeline and individual deployment actions for Amazon ECS, Amazon EC2, and AWS Lambda, along with support for manual approval actions, is the most suitable solution to meet the requirements of the continuous delivery pipeline.

B mentions using AWS CodeDeploy as the deploy provider, but it does not explicitly mention support for deploying to Amazon ECS, Amazon EC2, and AWS Lambda. AWS CodeDeploy primarily focuses on deploying applications to Amazon EC2 instances, and while it does have support for AWS Lambda, it might not be as straightforward to use for deploying to Amazon ECS.

upvoted 2 times

🗨️ 👤 **Aja1** 1 year, 7 months ago

i think B is correct

upvoted 1 times

🗨️ 👤 **habros** 1 year, 8 months ago

Selected Answer: B

You will need a deployment tool (CodeDeploy) for this. You cannot directly deploy via CodePipeline. Hence, B.

upvoted 6 times


A company has an application that runs on Amazon EC2 instances that are in an Auto Scaling group. When the application starts up, the application needs to process data from an Amazon S3 bucket before the application can start to serve requests. The size of the data that is stored in the S3 bucket is growing. When the Auto Scaling group adds new instances, the application now takes several minutes to download and process the data before the application can serve requests. The company must reduce the time that elapses before new EC2 instances are ready to serve requests.

Which solution is the MOST cost-effective way to reduce the application startup time?

- A. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Stopped state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
- B. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
- C. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Running state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
- D. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook and to place the new instance in the Standby state when the application is ready to serve requests.

Suggested Answer: C

Community vote distribution



haazybanj Highly Voted 1 year, 10 months ago

Selected Answer: A

Option A is the most cost-effective solution. By configuring a warm pool of EC2 instances in the Stopped state, the company can reduce the time it takes for new instances to be ready to serve requests. When the Auto Scaling group launches a new instance, it can attach the stopped EC2 instance from the warm pool. The instance can then be started up immediately, rather than having to wait for the data to be downloaded and processed. This reduces the overall startup time for the application.

Option C is also a solution that involves a warm pool of EC2 instances, but the instances are in the Running state. This means that they are already running and incurring costs, even though they are not currently serving requests. This is not a cost-effective solution.

upvoted 18 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: A

keywords: MOST cost-effective way to reduce the application startup time

upvoted 2 times

stoy123 11 months, 3 weeks ago

Selected Answer: C

C " The company must reduce the time that elapses before new EC2 instances are ready to serve requests."!!!!!!!!! this cannot happen with a stopped instance as it will still need to read the data from S3 upon startup,

upvoted 1 times

Jay_2pt0_1 10 months, 2 weeks ago

I thought this, as well, but A appears to be correct. See <https://aws.amazon.com/blogs/compute/scaling-your-applications-faster-with-ec2-auto-scaling-warm-pools/>

upvoted 1 times

Shasha1 1 year ago

A

for warm pool in the hibernated or stop status we will pay only for the attached EBS volume, therefore its much cost effective rather than running instance

upvoted 1 times

🗳️ 👤 **dzn** 1 year ago

Selected Answer: A

Warm Pool allows instances to be set to a stopped state after performing any process (e.g., running initialization scripts, warm-up tasks, etc.).
upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

A is correct: the question says <the application needs to process data from an Amazon S3 bucket before the application can start to serve requests> but <The size of the data that is stored in the S3 bucket is growing>. This means we should maintain a warm pool for EC2 so that they are always ready to process data (reduce the time that elapses before new EC2 instances are ready)

B and D: no mention of warmpool

C: If the instance is up and running, no need to configure warm pool

upvoted 1 times

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

Selected Answer: A

Answer is A, the question is cost-effective, and even with A you will have less wait time to download the S3 data, it will download the delta from the warm up process to ready to join to ASG

upvoted 2 times

🗳️ 👤 **Jaguaroooo** 1 year, 2 months ago

A&C are both good in terms of solutions, however, the caveat here is the "cost-effective" solution and that's why I agree with A.

<https://aws.amazon.com/blogs/compute/scaling-your-applications-faster-with-ec2-auto-scaling-warm-pools/>

upvoted 1 times

🗳️ 👤 **RVivek** 1 year, 6 months ago

Selected Answer: A

excerpt from the url: <https://aws.amazon.com/blogs/compute/scaling-your-applications-faster-with-ec2-auto-scaling-warm-pools/>

EC2 Auto Scaling Warm Pools works by launching a configured number of EC2 instances in the background, allowing any lengthy application initialization processes to run as necessary, and then stopping those instances until they are needed

upvoted 1 times

🗳️ 👤 **beanxyz** 1 year, 6 months ago

Selected Answer: A

A is the most cost-effective solution. Besides, when the warm EC2 was created, it already downloaded the contents from S3 so the next time when it started, it would just download any new files from S3. (e.g s3 sync)

upvoted 1 times

🗳️ 👤 **ixdb** 1 year, 6 months ago

C is right.

please carefully check the question:

The company must reduce the time that elapses before new EC2 instances are ready to serve requests.

When the application starts up. the application needs to process data from an Amazon S3 bucket before the application can start to serve requests.

upvoted 2 times

🗳️ 👤 **Chetantest07** 1 year, 7 months ago

Selected Answer: C

I understand the question is asking for the most cost-effective. keeping it stopped state is most cost efficient but it would not work because in the question it also states that, "When the application starts up. the application needs to process data" and to process that data takes time. If the Ec2 instance is stopped then started at the time of need, then again it will take time to process the data, right? so in this scenario, the EC2 instance need to be running.

upvoted 3 times

🗳️ 👤 **Suix** 1 year, 7 months ago

I think it should be C, as the A option would not be effective. Coming from the instance stop state the application will start up again and need to process the data from S3 bucket.

upvoted 2 times

🗳️ 👤 **Snape** 1 year, 8 months ago

Selected Answer: A

Warm pool with stopped state is most cost efficient option

upvoted 3 times

🗳️ 👤 **pepecastr0** 1 year, 8 months ago

Selected Answer: A

A - Keep it stopped until you need it to save money

upvoted 1 times

🗨️ 👤 **jqso234** 1 year, 11 months ago

Selected Answer: C

While A can also be a cost-effective solution, C is the MOST cost-effective solution because it utilizes Amazon S3 Transfer Acceleration, which is a feature that enables fast, easy, and secure transfers of files over the internet between Amazon S3 buckets and EC2 instances located in different regions or across the internet. By using S3 Transfer Acceleration, the data transfer speed can be increased significantly, which can reduce the time that elapses before new EC2 instances are ready to serve requests.

In contrast, A suggests using a larger instance size with more CPU and network capacity, which can be more expensive than the current instance size. Moreover, this approach may not be scalable in the long run since as the data in the S3 bucket continues to grow, the instance size may need to be further increased, which can incur more costs. Therefore, while A can also be a viable solution, C is the most cost-effective and scalable solution.

upvoted 1 times

🗨️ 👤 **ma_rio** 1 year, 11 months ago

Selected Answer: A

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

Keeping instances in a Stopped state is an effective way to minimize costs.

upvoted 4 times

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

Selected Answer: A

A for me to decrease costs

upvoted 2 times

A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.

The buildspec.yml file contains the following:

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```


The DevOps engineer has noticed that anybody with an AWS account is able to download the artifacts.

What steps should the DevOps engineer take to stop this?

- A. Modify the post_build command to use --acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.
- B. Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants read-only access.
- C. Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal "*".
- D. Modify the post_build command to remove --acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

Suggested Answer: D

Community vote distribution



haazybanj Highly Voted 1 year, 10 months ago

Selected Answer: D

D is correct

upvoted 13 times

beanxyz Highly Voted 1 year, 6 months ago

Selected Answer: A

--acl authenticated-read means any authenticated users can read the S3 bucket. We should remove it and configure the bucket policy to explicitly grant access

upvoted 5 times

beanxyz 1 year, 6 months ago

I mean D...

upvoted 6 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: D

"--acl authenticated-read" means any authenticated users can read the S3 bucket. We should remove it and configure the bucket policy to explicitly grant access

upvoted 3 times

zijo 10 months, 4 weeks ago

D is the answer

ACL-authenticated users: This refers to any user who has successfully authenticated with AWS credentials, including IAM users and federated users. It does not include anonymous users (public access).

It's generally recommended to use bucket policies for access control in S3 rather than ACLs. Bucket policies offer more granular control and better security practices. You can achieve "acl-authenticated reads" access using a bucket policy as well.


upvoted 2 times

  **dzn** 1 year ago

Selected Answer: D

`remove --acl authenticated-read` is required to fulfill the requirement.

upvoted 4 times

  **thanhnv142** 1 year, 1 month ago


B is correct: In the "buildspec.yml file", we see that there is "--acl authenticated-read". This allow all aws users who successfully authen to AWS can download the file. To restrict access, we need to modify ACL that only grant access to some specific users.

Note that we should not use bucket policy because it will affect all objects in the bucket (that is why it is called BUCKET policy). We only need to restrict access to an object, then ACL is the right choice.

A is incorrect: Use use --acl public-read means we allow all user to access the object

C and D: Use bucket policy, which is incorrect

upvoted 1 times

  **zolthar_z** 1 year, 3 months ago

Selected Answer: D

D is correct

upvoted 2 times

A company has developed a serverless web application that is hosted on AWS. The application consists of Amazon S3, Amazon API Gateway, several AWS Lambda functions, and an Amazon RDS for MySQL database. The company is using AWS CodeCommit to store the source code. The source code is a combination of AWS Serverless Application Model (AWS SAM) templates and Python code.


A security audit and penetration test reveal that user names and passwords for authentication to the database are hardcoded within CodeCommit repositories. A DevOps engineer must implement a solution to automatically detect and prevent hardcoded secrets.

What is the MOST secure solution that meets these requirements?

- A. Enable Amazon CodeGuru Profiler. Decorate the handler function with `@with_lambda_profiler()`. Manually review the recommendation report. Write the secret to AWS Systems Manager Parameter Store as a secure string. Update the SAM templates and the Python code to pull the secret from Parameter Store.
- B. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Manually check the code review for any recommendations. Choose the option to protect the secret. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
- C. Enable Amazon CodeGuru Profiler. Decorate the handler function with `@with_lambda_profiler()`. Manually review the recommendation report. Choose the option to protect the secret. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
- D. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Manually check the code review for any recommendations. Write the secret to AWS Systems Manager Parameter Store as a string. Update the SAM templates and the Python code to pull the secret from Parameter Store.

Suggested Answer: B -

Community vote distribution



haazybanj Highly Voted 1 year, 4 months ago

Selected Answer: B

B

The MOST secure solution that meets the requirement of automatically detecting and preventing hardcoded secrets is to use AWS CodeGuru Reviewer to check the code for any hardcoded secrets, and then update the SAM templates and Python code to retrieve the secrets from AWS Secrets Manager.

Option B is the correct answer. By associating the CodeCommit repository with Amazon CodeGuru Reviewer, the code can be checked for any hardcoded secrets during code reviews. When a hardcoded secret is detected, CodeGuru Reviewer will recommend updating the code to retrieve the secret from a secure storage service like AWS Secrets Manager. The DevOps engineer can choose the option to protect the secret and then update the SAM templates and Python code to retrieve the secret from AWS Secrets Manager instead of hardcoding it in the code.

upvoted 12 times

rhinozD Highly Voted 1 year, 2 months ago

Selected Answer: B

B is correct.

CodeGuru Reviewer for security problems.

Amazon CodeGuru Profiler is for performance.

upvoted 8 times

thanhnv142 Most Recent 7 months, 1 week ago

B is correct: <implement a solution to automatically detect and prevent hardcoded secrets> means we need CodeGuru reviewer to analyze the code and uncover hardcoded credentials.

A and C: no mention of CodeGuru reviewer

D: using System Manager Parameter store is a good method to avoid hardcoded credentials. However, the question requires <the MOST secure solution>, so we should use AWS secret manager (option B). It costs more than Para store, but more secure.

upvoted 1 times



a16a848 8 months, 2 weeks ago

Selected Answer: C

I'd say it's C, because the system to examine includes Python code and CodeGuru profiles for Python needs the decorator:

<https://docs.aws.amazon.com/codeguru/latest/profiler-ug/python-lambda.html>



upvoted 1 times

  **Aja1** 1 year, 1 month ago

option C

<https://docs.aws.amazon.com/codeguru/latest/profiler-ug/python-lambda-command-line.html>

upvoted 1 times

  **Aja1** 1 year, 1 month ago



Sorry B

Amazon CodeGuru Reviewer and Amazon CodeGuru Profiler are both tools that can be used to improve the quality and security of your code. However, they have different strengths and weaknesses.

CodeGuru Reviewer is a static code analysis tool that can be used to find potential defects in your code. It can scan your code for hardcoded secrets, security vulnerabilities, and other potential problems. CodeGuru Reviewer can also provide recommendations on how to fix the problems that it finds.

CodeGuru Profiler is a dynamic code analysis tool that can be used to understand how your code performs. It can track the performance of your code, identify bottlenecks, and suggest ways to improve performance. CodeGuru Profiler can also be used to find potential memory leaks and other performance problems.

upvoted 6 times

  **MarDog** 1 year, 2 months ago

Selected Answer: B

Definitely B.

upvoted 2 times

A company is using Amazon S3 buckets to store important documents. The company discovers that some S3 buckets are not encrypted. Currently, the company's IAM users can create new S3 buckets without encryption. The company is implementing a new requirement that all S3 buckets must be encrypted.


A DevOps engineer must implement a solution to ensure that server-side encryption is enabled on all existing S3 buckets and all new S3 buckets. The encryption must be enabled on new S3 buckets as soon as the S3 buckets are created. The default encryption type must be 256-bit Advanced Encryption Standard (AES-256).

Which solution will meet these requirements?

- A. Create an AWS Lambda function that is invoked periodically by an Amazon EventBridge scheduled rule. Program the Lambda function to scan all current S3 buckets for encryption status and to set AES-256 as the default encryption for any S3 bucket that does not have an encryption configuration.
- B. Set up and activate the s3-bucket-server-side-encryption-enabled AWS Config managed rule. Configure the rule to use the AWS-EnableS3BucketEncryption AWS Systems Manager Automation runbook as the remediation action. Manually run the re-evaluation process to ensure that existing S3 buckets are compliant.
- C. Create an AWS Lambda function that is invoked by an Amazon EventBridge event rule. Define the rule with an event pattern that matches the creation of new S3 buckets. Program the Lambda function to parse the EventBridge event, check the configuration of the S3 buckets from the event, and set AES-256 as the default encryption.
- D. Configure an IAM policy that denies the s3:CreateBucket action if the s3:x-amz-server-side-encryption condition key has a value that is not AES-256. Create an IAM group for all the company's IAM users. Associate the IAM policy with the IAM group.

Suggested Answer: D

Community vote distribution



paali Highly Voted 1 year, 9 months ago

B caters to both existing and new buckets.

C is triggered on when new bucket is created, existing buckets are not handled by the event.

upvoted 13 times

Zoe_zoe Highly Voted 1 year, 10 months ago

Selected Answer: B

B to me

upvoted 10 times

Gomer Most Recent 9 months ago

Selected Answer: C

I think neither "B" or "C" is complete solution. They both need to be done to deal with both existing and new buckets.

A carefull reading of the question doesn't preclude the need to do both.

However, the specific and emphasized criteria of enabling encryption "as soon as the S3 buckets are created" can only be done by "C" (event driven action)

I think this may be a trick question. I'm very confident they are defining an event driven action as part of the solution, and only "C" provides that.

B: (NO) "Manually run the re-evaluation process to ensure that existing S3 buckets are compliant."

Comment: Doesn't achieve "encryption must be enabled on new S3 buckets as soon as the S3 buckets are created."

upvoted 2 times

dzn 1 year ago

Selected Answer: B

`s3-bucket-server-side-encryption-enabled` checks if your Amazon S3 bucket either has the Amazon S3 default encryption enabled or that the Amazon S3 bucket policy explicitly denies put-object requests without server side encryption that uses AES-256 or AWS Key Management Service.

upvoted 1 times

thanhnv142 1 year, 1 month ago

A is correct: <implement a solution to ensure that server-side encryption is enabled on all existing S3 buckets and all new S3 buckets>. We can use lambda to configure all S3. Use Eventbridge to schedule-run lambda.

B: This option uses AWS config rule to activate AWS-EnableS3BucketEncryption AWS Systems Manager Automation runbook, which is incorrect. Remember that AWS config have no action and cannot trigger anything. It only collect data and report. Additionally, this option does not mention actions to new S3 bucket

C: <define the rule with an event pattern that matches the creation of new S3 buckets> means that this only affect newly-created bucket, not existing ones.

D: No mention of enforcing encryption on S3

Note: Should not use chatgpt for this exam, its answers are mostly wrong

upvoted 2 times

  **thanhnv142** 1 year ago

Correct: D

upvoted 1 times

  **davdan99** 1 year, 2 months ago

Selected Answer: B

Answer is B

<https://docs.aws.amazon.com/config/latest/developerguide/s3-bucket-server-side-encryption-enabled.html>

upvoted 5 times

  **Jaguaroooo** 1 year, 2 months ago

I would have chose B over D because aws config can do this with lambda.

upvoted 1 times

  **Jaguaroooo** 1 year, 2 months ago

A has automation. I didn't like B: because of this statement: Manually run the re-evaluation process to ensure that existing S3 buckets are compliant.

upvoted 1 times



  **Jamshif01** 1 year, 2 months ago

Amazon S3 Encrypts New Objects By Default

[https://aws.amazon.com/blogs/aws/amazon-s3-encrypts-new-objects-by-](https://aws.amazon.com/blogs/aws/amazon-s3-encrypts-new-objects-by-default/#:~:text=At%20AWS%2C%20security%20is%20the,specify%20a%20different%20encryption%20option.)

[default/#:~:text=At%20AWS%2C%20security%20is%20the,specify%20a%20different%20encryption%20option.](https://aws.amazon.com/blogs/aws/amazon-s3-encrypts-new-objects-by-default/#:~:text=At%20AWS%2C%20security%20is%20the,specify%20a%20different%20encryption%20option.)

upvoted 1 times



  **zenith_cloud** 1 year, 4 months ago

Selected Answer: B

B to me.

AWS Config can monitor resource compliance against desired configurations. The managed rule s3-bucket-server-side-encryption-enabled checks whether Amazon S3 buckets have server-side encryption enabled. The AWS Systems Manager Automation runbook, AWS-EnableS3BucketEncryption, can be used as a remediation action to enable default encryption. This solution would also work for new buckets as soon as they're created, making it an effective solution.

upvoted 1 times

  **rhinozD** 1 year, 8 months ago

Selected Answer: B

B is right.

Doable solution for new buckets as well as existing buckets.



upvoted 4 times

  **marcoforexam** 1 year, 10 months ago

Selected Answer: C

Option C meets the requirement of modifying the policy immediately after creating the bucket.

upvoted 1 times

  **rhinozD** 1 year, 8 months ago

What about existing buckets?

upvoted 1 times


A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons, users subscribing to this application are distributed across multiple Application Load Balancers (ALBs), each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage, and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs, Auto Scaling groups, and EC2 fleets.

Which architecture will meet these requirements with the LEAST amount of configuration?

- A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
- B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.
- C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
- D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB-Auto Scaling group pair.

Suggested Answer: C

Community vote distribution



rhinozD Highly Voted 1 year, 8 months ago

Selected Answer: C

You can just use one CodeDeploy application and multiple deployment groups in this case.
so C.
upvoted 12 times

steli0 Most Recent 3 months, 2 weeks ago

Selected Answer: C

I was about to vote B since the link from xdkonorek2 shows that one deployment can include up to 10 ELBs. Nevertheless the question says multiple instead of 10.
upvoted 1 times

trungtd 8 months ago

Selected Answer: C

Option B not feasible as it assumes a single deployment group can manage deployments across multiple ALBs and Auto Scaling groups simultaneously, which is not supported.
upvoted 1 times

xdkonorek2 8 months, 3 weeks ago

Selected Answer: B

<https://aws.amazon.com/about-aws/whats-new/2023/10/aws-codedeploy-multiple-load-balancers-amazon-ec2-applications/>
upvoted 2 times

xdkonorek2 10 months, 3 weeks ago

Selected Answer: B

B is the simplest :)

During creation of deployment group:

1. select "Amazon EC2 Auto Scaling groups"
 2. tip appears: "You can select up to 10 Amazon EC2 Auto Scaling groups to deploy your application revision to."
- upvoted 3 times

that1guy 10 months, 3 weeks ago

Also B for me, you can target multiple ASGs as part of one deployment:

https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_TargetInstances.html#CodeDeploy-Type-TargetInstances-

autoScalingGroups

upvoted 2 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

C is correct: <the application must trigger a simultaneous deployment> means deployment in parallel

B and D: no mention of deployment in parallel

A: <unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair> means there are multiple AWS CodeDeploy applications and deployment groups for each site, which is unnecessary

upvoted 3 times

🗨️ 👤 **Aja1** 1 year, 7 months ago

Option C

deployed in parallel to all ALB-Auto Scaling group pairs simultaneously. This means that the deployment process is efficient and fast, and all ALBs and Auto Scaling groups receive updates at the same time.

upvoted 1 times

🗨️ 👤 **devnv** 1 year, 10 months ago

C is the correct answer.

upvoted 3 times

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

Selected Answer: C

C is correct.

upvoted 3 times

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

Selected Answer: A

A

AWS CodePipeline can target multiple AWS CodeDeploy applications.

upvoted 1 times

A company is hosting a static website from an Amazon S3 bucket. The website is available to customers at example.com. The company uses an Amazon Route 53 weighted routing policy with a TTL of 1 day. The company has decided to replace the existing static website with a dynamic web application. The dynamic web application uses an Application Load Balancer (ALB) in front of a fleet of Amazon EC2 instances.

On the day of production launch to customers, the company creates an additional Route 53 weighted DNS record entry that points to the ALB with a weight of 255 and a TTL of 1 hour. Two days later, a DevOps engineer notices that the previous static website is displayed sometimes when customers navigate to example.com.

How can the DevOps engineer ensure that the company serves only dynamic content for example.com?

- A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.
- B. Update the weighted DNS record entry that points to the S3 bucket. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.
- C. Configure webpage redirect requests on the S3 bucket with a hostname that redirects to the ALB.
- D. Remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zone. Wait for DNS propagation to become complete.

Suggested Answer: D

Community vote distribution

D (82%)

B (18%)

🗳️ 👤 **2pk** Highly Voted 1 year, 4 months ago

Selected Answer: D

D is the answer

B wrong because

Route 53 initially considers only the nonzero weighted records, if any.

If all the records that have a weight greater than 0 are unhealthy, then Route 53 considers the zero-weighted records.

upvoted 9 times

🗳️ 👤 **steli0** Most Recent 3 months, 2 weeks ago

Selected Answer: D

B would be correct if instead of domain reset option (which doesn't exist) there was a TTL decrease or nothing.

upvoted 1 times

🗳️ 👤 **jamesf** 7 months, 1 week ago

Selected Answer: D

D is correct

Not B as

- Route 53 initially considers only the nonzero weighted records, if any.

- If all the records that have a weight greater than 0 are unhealthy, then Route 53 considers the zero-weighted records.

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-weighted.html>

Besides, in B, <domain reset option to propagate changes immediately>, i don't think DNS record will update immediately

upvoted 2 times

🗳️ 👤 **jojom19980** 1 year ago

Selected Answer: B

D is correct but I will go with B because is more save to the customers, some clients have the old record (TTL 1 day) so after 1 day I can confirm that all the clients have the the new DNS record so I can delete the record

upvoted 1 times

🗳️ 👤 **thanhv142** 1 year, 1 month ago

D is correct:

A: should not delete all objects from S3, they change nothing

C: should not do this, we have a more efficient method

B: < domain reset option to propagate changes immediately>: there is no such thing. DNS record will expire after TTL. Cannot force DNS resolvers to query for DNS record before TTL expire

upvoted 4 times

🗳️ 👤 **RVivek** 1 year, 6 months ago

Selected Answer: D

B- is incorrect , as gigi_devops has metioned setting 0 weight is not enough. Also reset domain option is not available.

upvoted 3 times

🗳️ 👤 **ixdb** 1 year, 6 months ago

D is right.

Just setting the weight to 0 does not ensure that traffic will not go to example.com.

upvoted 2 times

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

Selected Answer: D

Setting the weight to 0 is not enough. So C is the best answer. <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-weighted.html#:~:text=Si%20tous%20les%20enregistrements%20dont%20le%20poids%20est%20sup%C3%A9rieur%20%C3%A0%200%20ne%20sont%20pas>

upvoted 3 times

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

Setting the weight to 0 is not enough. So C is the best answer. <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-weighted.html#:~:text=Si%20tous%20les%20enregistrements%20dont%20le%20poids%20est%20sup%C3%A9rieur%20%C3%A0%200%20ne%20sont%20pas>

upvoted 1 times

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

Sorry I wanted to say "D"

upvoted 4 times

🗳️ 👤 **Aja1** 1 year, 7 months ago

option B is the most appropriate choice as it immediately redirects all traffic away from the S3 bucket and ensures that only the dynamic content from the ALB is served for example.com.

upvoted 1 times

🗳️ 👤 **Aja1** 1 year, 7 months ago

D.

We can remove the Weighted DNS as it is not necessary

upvoted 2 times

🗳️ 👤 **habros** 1 year, 8 months ago

Selected Answer: D

D. Also, do check the weight of the CNAME record of the ALB. It might be conflicting.

upvoted 1 times

🗳️ 👤 **Blueee** 1 year, 8 months ago

Selected Answer: D

agree with D

upvoted 1 times

🗳️ 👤 **MarDog** 1 year, 8 months ago

In reference to B, I don't think the "domain reset option" exists. So, it's D.

upvoted 4 times

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

Selected Answer: D

Option D.



upvoted 2 times

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

Selected Answer: D

To ensure that the company only serves dynamic content for example.com, the DevOps engineer should remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zone. This will immediately remove the static website from the DNS resolution pool. Solution B would only update the weight of the record entry, but it would still take 24 hours for the changes to propagate.

upvoted 4 times


  **rhinozD** 1 year, 8 months ago

Selected Answer: B

B

To disable routing to a resource, set Weight to 0

upvoted 2 times

  **nocinfra** 1 year, 9 months ago

Selected Answer: B

nocinfra 0 minutes ago Awaiting moderator approval

B for me , You can gradually change the balance by changing the weights. If you want to stop sending traffic to a resource, you can change the weight for that record to 0.

upvoted 2 times

A company is implementing AWS CodePipeline to automate its testing process. The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon EventBridge:

```
{
  "source": [
    "aws.codepipeline"
  ],
  "detail-type": [
    "CodePipeline Action Execution State Change"
  ],
  "detail": {
    "state": [
      "FAILED"
    ],
    "type": {
      "category": ["Approval"]
    }
  }
}
```

Which type of events will match this event pattern?

- A. Failed deploy and build actions across all the pipelines
- B. All rejected or failed approval actions across all the pipelines
- C. All the events across all pipelines
- D. Approval actions across all the pipelines

Suggested Answer: B

Community vote distribution

B (100%)

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

Selected Answer: B

Use this sample event pattern to capture all rejected or failed approval actions across all the pipelines.

<https://docs.aws.amazon.com/codepipeline/latest/userguide/detect-state-changes-cloudwatch-events.html>


upvoted 14 times

 **thanhnv142** Most Recent 7 months, 1 week ago

B is correct: <state:failed and category:approval> means failed approval

A, C and D: no mention of approval

upvoted 3 times

 **gdtypk** 1 year, 3 months ago

Selected Answer: B

https://docs.aws.amazon.com/ja_jp/codepipeline/latest/userguide/detect-state-changes-cloudwatch-events.html

upvoted 3 times

 **marcoforexam** 1 year, 4 months ago

Selected Answer: B

category: approval

upvoted 1 times

An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched, and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files in source control.

Which solution will accomplish this?

- A. In the CloudFormation template, add an AWS Config rule. Place the configuration file content in the rule's InputParameters property, and set the Scope property to the EC2 Auto Scaling group. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- B. In the CloudFormation template, add an EC2 launch template resource. Place the configuration file content in the launch template. Configure the cfn-init script to run when the instance is launched, and configure the cfn-hup script to poll for updates to the configuration.
- C. In the CloudFormation template, add an EC2 launch template resource. Place the configuration file content in the launch template. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- D. In the CloudFormation template, add CloudFormation init metadata. Place the configuration file content in the metadata. Configure the cfn-init script to run when the instance is launched, and configure the cfn-hup script to poll for updates to the configuration.

Suggested Answer: B

Community vote distribution

D (79%)

B (21%)

4555894 Highly Voted 1 year ago

Selected Answer: D

Use the AWS::CloudFormation::Init type to include metadata on an Amazon EC2 instance for the cfn-init helper script. If your template calls the cfn-init script, the script looks for resource metadata rooted in the AWS::CloudFormation::Init metadata key. Reference:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html>

upvoted 7 times

steli0 Most Recent 3 months, 2 weeks ago

Selected Answer: D

It's D

upvoted 1 times

jamesf 7 months, 2 weeks ago

Selected Answer: D

Require CloudFormation init for cfn-init and cfn-hup

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-hup.html>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html>

upvoted 2 times

[Removed] 1 year ago

Selected Answer: D

In this scenario, CloudFormation init metadata is the most suitable approach for ensuring that instances launched by the Auto Scaling group have the latest configuration file.

upvoted 2 times

thanhnv142 1 year, 1 month ago

D is correct: <The instances are created and maintained with AWS CloudFormation> means we will only use ACF to satisfy the requirements of this question. <changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated> means we need cfn-init, which is a daemon that check for updates and update the changes

A and C: no mention of cfn-init

B: no mention of CloudFormation init. we need CloudFormation init because cfn-init is specified in CloudFormation init key.

upvoted 3 times

🗳️ 👤 **khchan123** 1 year, 1 month ago

Selected Answer: D

D. cfn-hup poll for cloudformation metadata. B is wrong because putting the config content in launch template instead of metadata, where cfn-hub is not able to poll.

upvoted 4 times

🗳️ 👤 **a54b16f** 1 year, 1 month ago

Selected Answer: D

cfn-init is defined inside AWS::CloudFormation::Init

upvoted 2 times

🗳️ 👤 **yuliaqwerty** 1 year, 2 months ago

I vote for B

upvoted 1 times

🗳️ 👤 **Jaguaroooo** 1 year, 2 months ago

But what happened to the aspect of using source control?

upvoted 4 times

🗳️ 👤 **a16a848** 1 year, 2 months ago

Selected Answer: B

Google Bart says it is B.

By using an EC2 launch template resource, the configuration file will be installed and configured on all instances when they are launched. The cfn-init script will also poll for updates to the configuration, so that all instances will have the latest configuration file as soon as it is updated.

In addition, the solution will comply with company policy by storing the configuration file in source control along with the AWS infrastructure configuration files. This will ensure that changes to the configuration file are tracked and managed in a consistent way.

Option C: Using an AWS Systems Manager Resource Data Sync resource alone is not enough to ensure that all instances have the latest configuration file. The cfn-init script is needed to install and configure the configuration file on the instance, and the cfn-hup script is needed to poll for updates to the configuration.

upvoted 2 times

🗳️ 👤 **svjl** 1 year, 3 months ago

Selected B:

You can have cfn on Launch Config and Launch Template

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-hup.html>

<https://stackoverflow.com/questions/54691327/cfn-init-for-cloudformation-launchtemplate>

upvoted 3 times

🗳️ 👤 **HugoFM** 1 year, 3 months ago

Selected Answer: B

I choos B because the config file must be maintained along teh infrastructure configuration files in source control

upvoted 1 times

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

Selected Answer: B

B and D are similar. I will go for B, because D doesn't involve EC2 launch templates

upvoted 1 times

🗳️ 👤 **RVivek** 1 year, 6 months ago

Selected Answer: D

cfn-init and cfn-hup are used to update metadata. <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-hup.html>

upvoted 2 times

🗳️ 👤 **Aja1** 1 year, 7 months ago

option C provides a reliable and scalable solution to manage the configuration file for the application running on the EC2 instances, while also adhering to company policies regarding source control for configuration files.


upvoted 1 times

🗳️ 👤 **Aja1** 1 year, 7 months ago

D

cfn-hup is a daemon that detects changes in resource metadata and runs user-specified actions when a change is detected. This allows you to automatically update the configuration of your Amazon EC2 instances when you make changes to your AWS CloudFormation stacks.

upvoted 4 times

  **rhinozD** 1 year, 8 months ago



Selected Answer: D

D

Metadata:

"AWS::CloudFormation::Init":

upvoted 4 times

  **2pk** 1 year, 10 months ago

Selected Answer: B

By using an EC2 launch template, you can include the configuration file content directly in the template. The cfn-init script can be configured to run when the instance is launched.

as CloudFormation init metadata is more suitable for configuring instances during stack creation rather than for dynamically updating configuration files.

upvoted 3 times

A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket. Logs are rarely accessed after 90 days and must be retained for 10 years.

Which combination of steps should a DevOps engineer take to meet these requirements? (Choose two.)

- A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.
- B. Configure a CloudWatch Logs subscription filter to use Amazon Kinesis Data Firehose to stream all logs to an S3 bucket.
- C. Configure a CloudWatch Logs subscription filter to stream all logs to an S3 bucket.
- D. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier after 90 days and to expire logs after 3.650 days.
- E. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3.650 days.

Suggested Answer: BD

Community vote distribution

BD (84%)

CD (16%)

🗳️ 👤 **2pk** Highly Voted 1 year, 10 months ago

Selected Answer: BD

Amazon Kinesis Data Firehose simplifies the process of loading streaming data into S3 and provides automatic scaling, buffering, and retries.
upvoted 6 times

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

Selected Answer: BD

B - keywords: continue stream but not one time task
D - keywords: S3 Glacier
upvoted 1 times

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

Selected Answer: BD

vote B and D.
I initially thought the C is better than B, because Amazon Kinesis Data Firehose is primarily used for time-sensitive tasks, which is not suitable for this case, But when I read the C. I found the Directly streaming logs from cloudwatch log to s3 is not a feature provided by Cloudwatch.
So, I will go with B and D.
upvoted 3 times

🗳️ 👤 **Gomer** 9 months ago

Selected Answer: BD

You can absolutely directly "export log data from your log groups to an Amazon S3 bucket"
However, this is a one time export, and NOT an ongoing stream.
If you want to stream continuously you have to use "subscription filter with Kinesis Data Streams, Lambda, or Firehose."
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/S3Export.html>
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>
<https://dev.to/aws-builders/automate-export-of-cloudwatch-logs-to-s3-bucket-using-lambda-with-eventbridge-trigger-2ieg>
upvoted 3 times

🗳️ 👤 **zijo** 9 months, 4 weeks ago

Looks like creating subscription filters in AWS cloudwatch logs, there are only limited destination options. There is no S3 as a direct destination. You have to either create Elasticsearch or Kinesis or Kinesis Firehose or Lambda subscription filters. Given the choices we have, we need to pick B & D
upvoted 1 times

🗳️ 👤 **Jay_2pt0_1** 10 months, 1 week ago

Selected Answer: CD

C & D for the reasons that thanhnv142 mentioned.
upvoted 1 times

🗳️ 👤 **vn_thanhung** 10 months ago

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>

Pls check link. You can use a subscription filter with Kinesis Data Streams, Lambda, or Firehose. Logs that are sent to a receiving service through a subscription filter are base64 encoded and compressed with the gzip format. correct is B and D

upvoted 2 times

🗳️ 👤 **Heyang** 1 year ago

CD The question does not mention trying to switch to S3 in real time. C is more cost-effective.

https://docs.aws.amazon.com/zh_cn/AmazonCloudWatch/latest/logs/S3ExportTasksConsole.html

upvoted 2 times

🗳️ 👤 **dzn** 1 year ago

Selected Answer: BD

Amazon S3 Glacier is a secure, durable, and very low-cost cloud storage service that can be used for data archiving and long-term backup.

upvoted 3 times

🗳️ 👤 **jojom19980** 1 year ago

Selected Answer: BD

You can use a subscription filter with Kinesis Data Streams, Lambda, or Firehose.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>

upvoted 2 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

Selected Answer: CD

C and D is correct: <archive the logs to an Amazon S3 bucket> means we need to transport logs from CloudWatch Logs to S3. CloudWatch Logs can directly transport log data to S3. Logs are rarely accessed after 90 days means we need S3 bucket lifecycle policy

upvoted 4 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

A: AWS Glue is used primarily to integrate data from multiple data sources (up to 70) for data analysis. Of course it works well with one data source only (CW logs in this case). But it costs a lot of money and using it with only one data source is a waste of corporate budget. Should not use this.

B: Amazon Kinesis Data Firehose is primarily used for time-sensitive tasks, such as video streaming. It is very powerful that it can handle data in near realtime. However, this premium feature comes with a big expense. We only need to archive data, not video streaming it.

E: we need to transition it to S3 Glacier not Reduced Redundancy after 90 days

upvoted 1 times

🗳️ 👤 **habros** 1 year, 8 months ago

Selected Answer: BD

B to shift logs out using Kinesis Firehose to S3. Then D to set S3 bucket storage class to Glacier Flexible.

upvoted 4 times

🗳️ 👤 **YXXt55** 1 year, 8 months ago

Selected Answer: BD

C would make sense, but subscription filters don't go to S3 directly

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>

upvoted 4 times

🗳️ 👤 **Aja1** 1 year, 7 months ago

Option C is incorrect because streaming all logs to an S3 bucket is not a good solution for archiving logs

upvoted 2 times

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

Selected Answer: BD

BD is right

upvoted 1 times

🗳️ 👤 **OrganizedChaos25** 1 year, 9 months ago

Selected Answer: BD

upvoted 1 times

A company is developing a new application. The application uses AWS Lambda functions for its compute tier. The company must use a canary deployment for any changes to the Lambda functions. Automated rollback must occur if any failures are reported.


The company's DevOps team needs to create the infrastructure as code (IaC) and the CI/CD pipeline for this solution.

Which combination of steps will meet these requirements? (Choose three.)

- A. Create an AWS CloudFormation template for the application. Define each Lambda function in the template by using the `AWS::Lambda::Function` resource type. In the template, include a version for the Lambda function by using the `AWS::Lambda::Version` resource type. Declare the `CodeSha256` property. Configure an `AWS::Lambda::Alias` resource that references the latest version of the Lambda function.
- B. Create an AWS Serverless Application Model (AWS SAM) template for the application. Define each Lambda function in the template by using the `AWS::Serverless::Function` resource type. For each function, include configurations for the `AutoPublishAlias` property and the `DeploymentPreference` property. Configure the deployment configuration type to `LambdaCanary10Percent10Minutes`.
- C. Create an AWS CodeCommit repository. Create an AWS CodePipeline pipeline. Use the CodeCommit repository in a new source stage that starts the pipeline. Create an AWS CodeBuild project to deploy the AWS Serverless Application Model (AWS SAM) template. Upload the template and source code to the CodeCommit repository. In the CodeCommit repository, create a `buildspec.yml` file that includes the commands to build and deploy the SAM application.
- D. Create an AWS CodeCommit repository. Create an AWS CodePipeline pipeline. Use the CodeCommit repository in a new source stage that starts the pipeline. Create an AWS CodeDeploy deployment group that is configured for canary deployments with a `DeploymentPreference` type of `Canary10Percent10Minutes`. Upload the AWS CloudFormation template and source code to the CodeCommit repository. In the CodeCommit repository, create an `appspect.yml` file that includes the commands to deploy the CloudFormation template.
- E. Create an Amazon CloudWatch composite alarm for all the Lambda functions. Configure an evaluation period and dimensions for Lambda. Configure the alarm to enter the ALARM state if any errors are detected or if there is insufficient data.
- F. Create an Amazon CloudWatch alarm for each Lambda function. Configure the alarms to enter the ALARM state if any errors are detected. Configure an evaluation period, dimensions for each Lambda function and version, and the namespace as `AWS/Lambda` on the Errors metric.

Suggested Answer: BCF

Community vote distribution



thanhnv142 7 months, 1 week ago

Selected Answer: BCF

BCF is my choice
upvoted 2 times

thanhnv142 7 months, 1 week ago

BCF are correct:

A is not correct: <needs to create the infrastructure as code (IaC)> means we prefer AWS SAM over ACF. ACF is used to deploy AWS instances, not for IaC

D is wrong: no mention of AWS SAM

E is wrong: <Amazon CloudWatch composite alarm for all the Lambda functions>, but we need alarm for each lambda func, not one alarm for all of them

upvoted 2 times

sarlos 8 months, 2 weeks ago

it should be BDF because code deploy can be configured for canary

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>

upvoted 2 times

HugoFM 9 months, 2 weeks ago

Selected Answer: BCF

BCF, E is not correct you need to monitor each lambda to do a rollback of a particular deploy

upvoted 3 times

🗨️ **AzureDP900** 9 months, 2 weeks ago

BCF is right

upvoted 1 times

🗨️ **r1f** 10 months, 2 weeks ago

Answer is BCF.

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>

upvoted 1 times

🗨️ **sivre** 10 months, 3 weeks ago

Can someone please explain why not A and D? seem the same of B C without using SAM

upvoted 3 times

🗨️ **RVivek** 12 months ago

Selected Answer: BCF

A is wrong beacuse of AWS::Lambda::Function

B - Can work

C: is correct

D: SAM or Lambda deployment in Codedeploy cannot be canary deployment. canaray deployment should be included in the Lambda code as mentioned in option B.

E: Composite Alram is not required. if any Lambda fails , it should generate alarm

F: works

Basically select B from AB which is for lambda coding, Select C from CD for deploying and F from EF for monitoring and alerting

upvoted 2 times

🗨️ **hotblooded** 7 months, 2 weeks ago

Why we cannot use code deploy for canary there are already few deployment percentage for codedeploy BDF is correct

upvoted 1 times

🗨️ **habros** 1 year, 2 months ago

Selected Answer: BCF

Leaning to BCF. Lambda errors are standard although BCE is possible too.

For server less it is giveaway question. Stick with SAM if possible

upvoted 2 times

🗨️ **habros** 1 year, 2 months ago

I'll stick with BCF still. Composite alarms does not apply in this context.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/Create_Composite_Alarm.html

upvoted 2 times

🗨️ **sb333** 1 year, 2 months ago

Selected Answer: BCF

BCF is correct.

upvoted 2 times

🗨️ **Blueee** 1 year, 2 months ago

Selected Answer: BCF

BC and F

upvoted 3 times

🗨️ **Manny20** 1 year, 2 months ago

Composite Alarm requires underlying metric alarms which requires one CloudWatch alarm for each lambda functions and then tie them back to a composite alarm. So BC and F makes sense.

upvoted 3 times

🗨️ **FunkyFresco** 1 year, 2 months ago

Selected Answer: BCE

I think BCE makes more sense.

upvoted 1 times

🗨️ **ducluanxutrieu** 1 year, 2 months ago

Selected Answer: BCE

F creates an Amazon CloudWatch alarm for each Lambda function. However, it is not necessary to create an alarm for each Lambda function. A single composite alarm can be used to monitor all the Lambda functions.

upvoted 1 times

🗨️ 👤 **sb333** 1 year, 2 months ago

The issue with Answer E is that the alarm will also trigger on "insufficient data", which is not what you want. Answer F is correct.

upvoted 1 times

🗨️ 👤 **OrganizedChaos25** 1 year, 3 months ago

BCF are correct

upvoted 1 times

🗨️ 👤 **2pk** 1 year, 3 months ago

Cannot deploy the canary deployment in a pipeline for lambda creation, it has to be created in lambda resource file.

upvoted 1 times

🗨️ 👤 **2pk** 1 year, 3 months ago

Selected Answer: BCF

BCF correct

upvoted 4 times


A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status, and code was not deployed in the instances associated with the deployment group.

What are valid reasons for this failure? (Choose two.)

- A. The networking configuration does not allow the EC2 instances to reach the internet via a NAT gateway or internet gateway, and the CodeDeploy endpoint cannot be reached.
- B. The IAM user who triggered the application deployment does not have permission to interact with the CodeDeploy endpoint.
- C. The target EC2 instances were not properly registered with the CodeDeploy endpoint.
- D. An instance profile with proper permissions was not attached to the target EC2 instances.
- E. The appspec.yml file was not included in the application revision.

Suggested Answer: AD

Community vote distribution



haazybanj Highly Voted 1 year, 7 months ago

Selected Answer: AD

A.

Explanation: For CodeDeploy to work, the EC2 instances need to reach the CodeDeploy endpoint to download the deployment artifacts. If the networking configuration of the EC2 instances does not allow them to access the internet via a NAT gateway or internet gateway, they won't be able to reach the CodeDeploy endpoint, leading to deployment failure.

D

Explanation: When EC2 instances are part of a CodeDeploy deployment group, they need to have an associated IAM instance profile with the necessary permissions to interact with CodeDeploy and download the deployment artifacts. If the instance profile with proper permissions is not attached to the target EC2 instances, the deployment will fail as the instances won't have the required permissions to complete the deployment process.

upvoted 11 times

rhinozD Highly Voted 1 year, 8 months ago

Selected Answer: AD

AD

<https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html>

Search with: Troubleshooting all lifecycle events skipped errors

upvoted 7 times

bnagaraja9099 1 year, 2 months ago

C is correct.

the first reason for skipped errors on the link.

The CodeDeploy agent might not be installed or running on the instance. To determine if the CodeDeploy agent is running:

upvoted 1 times

sejar 1 year ago

No registration required, once agent is installed it should be sufficient. However permissions and network connectivity to S3 or code deploy would be must. Since that takes priority, A&D should be right.

upvoted 1 times

YucelFuat Most Recent 6 months ago

My question is "skipped" situation doesn't sound like a network error. There is no 4xx error or fail status

upvoted 1 times

zijo 9 months, 4 weeks ago

The user needs to create a service role and attach the `AWSCodeDeployRole` policy to it to grant the correct permissions for CodeDeploy to access EC2 instances. The role chosen should allow access to start and stop EC2 instances.

If the IAM role used by CodeDeploy doesn't have the necessary permissions to access the deployment artifacts or interact with the EC2 instances, the deployment may be skipped.

So it is not the IAM permissions of the user invoking the CodeDeploy.

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

Selected Answer: AD

A and D are correct: the deployment process might be skipped because of codedeploy agent

A: no connection means skipped deployment

D: insufficient permission means skipped deployment

upvoted 1 times

🗳️ 👤 **khchan123** 1 year, 1 month ago

Selected Answer: AD

A and D. See <https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-skipped-lifecycle-events>

upvoted 1 times

🗳️ 👤 **3a29cc4** 1 year, 2 months ago

Do you really have to have internet connectivity to use CodeDeploy? Why not use VPC endpoint in such cases? I go for CD.

upvoted 1 times

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

Selected Answer: CD

Some of the other options could cause a deployment to fail, but not specifically result in a "Skipped" status:

A Networking issues may prevent the deployment from reaching instances, but this would likely cause the deployment to fail, not be skipped.

B Lack of permissions for the IAM user would cause the deployment job itself to fail authorization.

E Missing `appspec.yml` would cause validation errors prior to the deployment attempt.

Anyone has different views?

upvoted 2 times

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

oh yeah;

why;

C -> CodeDeploy needs to be able to communicate with the instances in order to deploy revisions to them. If the instances are not registered, CodeDeploy will skip deploying to them.

D -> i think everyone know that point. i guess dont need explaintion.

Peace :)

upvoted 1 times

🗳️ 👤 **RVivek** 1 year, 6 months ago

Selected Answer: AD

I agree with rhinozD

upvoted 1 times

🗳️ 👤 **Blueee** 1 year, 8 months ago

Selected Answer: AD

AD is correct

upvoted 1 times

🗳️ 👤 **devnv** 1 year, 10 months ago

Its AD

upvoted 1 times

🗳️ 👤 **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: AD

AD is correct

upvoted 1 times

A company has a guideline that every Amazon EC2 instance must be launched from an AMI that the company's security team produces. Every month, the security team sends an email message with the latest approved AMIs to all the development teams.

The development teams use AWS CloudFormation to deploy their applications. When developers launch a new service, they have to search their email for the latest AMIs that the security department sent. A DevOps engineer wants to automate the process that the security team uses to provide the AMI IDs to the development teams.

What is the MOST scalable solution that meets these requirements?

- A. Direct the security team to use CloudFormation to create new versions of the AMIs and to list the AMI ARNs in an encrypted Amazon S3 object as part of the stack's Outputs section. Instruct the developers to use a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.
- B. Direct the security team to use a CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs and places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline output. Instruct the developers to use a cross-stack reference within their own CloudFormation template to obtain the S3 object location and the most recent AMI ARNs.
- C. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to place the AMI ARNs as parameters in AWS Systems Manager Parameter Store. Instruct the developers to specify a parameter of type SSM in their CloudFormation stack to obtain the most recent AMI ARNs from Parameter Store.
- D. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to create an Amazon Simple Notification Service (Amazon SNS) topic so that every development team can receive notifications. When the development teams receive a notification, instruct them to write an AWS Lambda function that will update their CloudFormation stack with the most recent AMI ARNs.

Suggested Answer: C

Community vote distribution

C (100%)

 **thanhv142** Highly Voted 1 year, 1 month ago

Selected Answer: C

C is correct: <automate the process that the security team uses to provide the AMI IDs to the development teams> and <MOST scalable solution> means we need a pipeline (image builder) to build AMI and to automate sharing

A and B: no mention of EC2 Image builder, which is better than codepipeline in building Ec2 image

D: They have to do this manually

upvoted 6 times

 **ad3fdb1** Most Recent 3 months, 3 weeks ago


A question to answer of option C - is it able to update the System Manager Parameter Store automatically? Option A seems able to do it automatically, right?

upvoted 1 times

 **yuliaqwerty** 1 year, 2 months ago

C is the best option


upvoted 2 times

 **rlf** 1 year, 4 months ago

Answer is C.

<https://aws.amazon.com/ko/blogs/compute/tracking-the-latest-server-images-in-amazon-ec2-image-builder-pipelines/>

upvoted 2 times

 **habros** 1 year, 8 months ago

Selected Answer: C

Use SSM Parameter Store or Secret Manager as the lookup K/V store for all the related AMIs. ANother way is also for security team to constantly update and share the images cross-account and grant them KMS keys to the encrypted AMIs. (not in question)

upvoted 2 times

 **devnv** 1 year, 10 months ago

C is correct

upvoted 2 times

  **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: C

C make more sense

upvoted 4 times

An application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). A DevOps engineer is using AWS CodeDeploy to release a new version. The deployment fails during the AllowTraffic lifecycle event, but a cause for the failure is not indicated in the deployment logs.

What would cause this?

- A. The appspec.yml file contains an invalid script that runs in the AllowTraffic lifecycle hook.
- B. The user who initiated the deployment does not have the necessary permissions to interact with the ALB.
- C. The health checks specified for the ALB target group are misconfigured.
- D. The CodeDeploy agent was not installed in the EC2 instances that are part of the ALB target group.

Suggested Answer: C

Community vote distribution

C (100%)


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

Selected Answer: C

C is the answer

refer this: <https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-deployments-allowtraffic-no-logs>

upvoted 17 times

 **thanhnv142** Highly Voted 7 months, 1 week ago

Selected Answer: C

C is correct: <deployment fails during the AllowTraffic lifecycle event> means there are problems with ALB.

A: no mention of the ALB

B: The user who init the deployment does not need necessary permission

D: If agent was not installed, it would fail from the start

upvoted 5 times

 **OrganizedChaos25** Most Recent 1 year, 3 months ago

C is the answer

upvoted 2 times

 **devnv** 1 year, 3 months ago

C is the correct answer

upvoted 1 times

A company has 20 service teams. Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192.168.0.0/22 CIDR block. The company manages the AWS accounts with AWS Organizations.

Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each other across the public internet. The company's security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.


A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team.

Which solution will meet these requirements?

- A. Create a new AWS account in AWS Organizations. Create a VPC in this account, and use AWS Resource Access Manager to share the private subnets of this VPC with the organization. Instruct the service teams to launch a new Network Load Balancer (NLB) and EC2 instances that use the shared private subnets. Use the NLB DNS names for communication between microservices.
- B. Create a Network Load Balancer (NLB) in each of the microservice VPCs. Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs. Create subscriptions to each VPC endpoint in each of the other AWS accounts. Use the VPC endpoint DNS names for communication between microservices.
- C. Create a Network Load Balancer (NLB) in each of the microservice VPCs. Create VPC peering connections between each of the microservice VPCs. Update the route tables for each VPC to use the peering links. Use the NLB DNS names for communication between microservices.
- D. Create a new AWS account in AWS Organizations. Create a transit gateway in this account, and use AWS Resource Access Manager to share the transit gateway with the organization. In each of the microservice VPCs, create a transit gateway attachment to the shared transit gateway. Update the route tables of each VPC to use the transit gateway. Create a Network Load Balancer (NLB) in each of the microservice VPCs. Use the NLB DNS names for communication between microservices.

Suggested Answer: B

Community vote distribution



Blueee Highly Voted 1 year, 8 months ago

Selected Answer: B

B is correct because all 20 services team in different separate AWS accounts are using the same CIDR block, which means they are overlapping CIDR.

D state that to update the route tables of each VPC to use the transit gateway but they are all having the same CIDR block so this cannot proceed, as shared by Arnaud92 link the pre-requisite of using the transit gateway is "No-overlapping CIDR block between VPCs."

upvoted 8 times

Saudis Most Recent 3 months, 4 weeks ago

Selected Answer: B

PrivateLink = HTTPS connection

upvoted 1 times

zijo 9 months, 3 weeks ago

B is the answer

When VPCs have overlapping CIDR blocks, AWS PrivateLink still ensures secure and private connectivity by using Interface Endpoints (ENIs) and Network Load Balancers (NLBs) to route traffic, bypassing the need for direct IP routing between the VPCs.

upvoted 2 times

thanhnv142 1 year, 1 month ago

B is correct: <all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet> means privatelink

A and C: no mention of privatelink

D: Using transite gateway. But this solution need IP to route traffic and cannot be used for overlapped VPC CIDR block (every team uses 192.168.0.0/22)

upvoted 2 times

🗨️ **zolphar_z** 1 year, 3 months ago

Selected Answer: B

Answer is B, Transit gateway can't route overlapping networks, the solution for this is privatelink:

<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-privatelink.html>

upvoted 4 times

🗨️ **RVivek** 1 year, 6 months ago

Selected Answer: D

Thanks to rhinozD. Please check the side by side comparison at the bottom of this page <https://tomgregory.com/cross-account-vpc-access-in-aws>

upvoted 2 times

🗨️ **[Removed]** 1 year, 5 months ago

In that same document you shared it says:

No-overlapping CIDR block between VPCs possible for Transit Gateway.

So it cannot be D.

upvoted 5 times

🗨️ **ixdb** 1 year, 6 months ago

B is right.,

upvoted 2 times

🗨️ **Just_Ninja** 1 year, 8 months ago

Selected Answer: B

B. is the right Solution!

Due to AWS's Transit Gateway not supporting same CIDRs (<https://aws.amazon.com/transit-gateway/faqs>), the most viable solution is the deployment of a Network Load Balancer (NLB) in each VPC. However, it's crucial to note that NLB operates similar to a NAT Gateway, allowing only incoming requests. After an incoming request is accepted, the NLB can then provide a response.

upvoted 3 times

🗨️ **SVGoogle89** 1 year, 8 months ago

AWS Transit Gateway doesn't support routing between Amazon VPCs with identical CIDRs. If you attach a new Amazon VPC that has a CIDR which is identical to an already attached Amazon VPC, AWS Transit Gateway will not propagate the new Amazon VPC route into the AWS Transit Gateway route table.

upvoted 1 times

🗨️ **habros** 1 year, 8 months ago

I'll lean towards B. For D, transit gateway is really expensive and does get the job done. There is also a need for NAT gateway as by default all AWS API traffic passes through the public internet. Hence, PrivateLink endpoints are for.

upvoted 1 times

🗨️ **FunkyFresco** 1 year, 8 months ago

Selected Answer: D

I go with option D. It makes more sense to me.

upvoted 1 times

🗨️ **allen_devops** 1 year, 8 months ago

I think the correct answer is B. Please note all service team is using the same cidr block for their vpc. It's impossible to add them in the same network mesh using vpc peering and transit gateway.

upvoted 3 times

🗨️ **Arnaud92** 1 year, 9 months ago

Selected Answer: D

see <https://tomgregory.com/cross-account-vpc-access-in-aws/> , Option 3

The use of a central hub reduce the complexity for 20 accounts

need an additional account to avoid cidr block collision, in the link they put the transit gateway in one of existing account

upvoted 2 times

🗨️ **rhinozD** 1 year, 8 months ago

Please read the "Side-by-side comparison" part at the end of the post.

D is wrong.

B is correct.

upvoted 2 times

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

Answer is D.

Option B is incorrect because it requires creating a Network Load Balancer in each of the microservice VPCs and using AWS PrivateLink to create VPC endpoints. This would result in a lot of configuration changes for each service team and increased complexity.

upvoted 1 times

🗨️ 👤 **devnv** 1 year, 10 months ago

B is the right answer

upvoted 2 times

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

Selected Answer: B

B is correct

upvoted 4 times

🗨️ 👤 **PhuocT** 1 year, 10 months ago

Option D is correct to me.

upvoted 3 times

An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received.

What are the possible causes for this error? (Choose two.)

- A. The S3 bucket default encryption is enabled.
- B. There is an error in the S3 bucket policy.
- C. The object has been moved to S3 Glacier.
- D. There is an error in the IAM role configuration.
- E. S3 Versioning is enabled.

Suggested Answer: BD

Community vote distribution

BD (100%)

🗳️ 👤 **juliaqwerty** 8 months ago

I think B and D
upvoted 3 times

🗳️ 👤 **Jamshif01** 8 months, 2 weeks ago

ACCESS DENIED - you got it
upvoted 2 times

🗳️ 👤 **RVivek** 12 months ago

Selected Answer: BD

IMHO it is BD
upvoted 3 times

🗳️ 👤 **vhernan** 1 year, 1 month ago

Selected Answer: BD

BD
Not an error though. Misconfiguration.
upvoted 4 times

🗳️ 👤 **FunkyFresco** 1 year, 2 months ago

Selected Answer: BD

B and D for sure.
upvoted 2 times

🗳️ 👤 **OrganizedChaos25** 1 year, 3 months ago

BD are the answers I got
upvoted 1 times

🗳️ 👤 **devnv** 1 year, 3 months ago

BD are correct
upvoted 1 times

A company wants to use a grid system for a proprietary enterprise in-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes, an `/etc/cluster/nodes.config` file must be updated, listing the IP addresses of the current node members of that cluster.

The company wants to automate the task of adding new nodes to a cluster.

What can a DevOps engineer do to meet these requirements?

- A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster. Create a Chef recipe that populates the content of the `/etc/cluster/nodes.config` file and restarts the service by using the current members of the layer. Assign that recipe to the Configure lifecycle event.
- B. Put the file `nodes.config` in version control. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for the cluster nodes. When adding a new node to the cluster, update the file with all tagged instances, and make a commit in version control. Deploy the new file and restart the services.
- C. Create an Amazon S3 bucket and upload a version of the `/etc/cluster/nodes.config` file. Create a crontab script that will poll for that S3 file and download it frequently. Use a process manager, such as Monit or systemd, to restart the cluster services when it detects that the new file was modified. When adding a node to the cluster, edit the file's most recent members. Upload the new file to the S3 bucket.
- D. Create a user data script that lists all members of the current security group of the cluster and automatically updates the `/etc/cluster/nodes.config` file whenever a new instance is added to the cluster.

Suggested Answer: A

Community vote distribution

A (100%)

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

Selected Answer: A

I'll use config management tool as well. In this case Opsworks (Chef/Puppet).
upvoted 6 times

 **Exto1124** 7 months, 2 weeks ago

But how the files content (get actual nodes list) is updated in that case?
upvoted 1 times

 **thanhnv142** Highly Voted 1 year, 1 month ago

Selected Answer: A


A is correct: <wants to use a grid system> means opswork stacks
B, C and D: no mention of opswork stack
upvoted 5 times

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


Selected Answer: A

AWS OpsWorks services have reached end of life and have been disabled for both new and existing customers. Will this question surface in the exam?

<https://aws.amazon.com/blogs/mt/migrate-your-aws-opsworks-stacks-to-aws-systems-manager/>
upvoted 2 times

 **hayjaykay** 4 months, 2 weeks ago

D.
This approach ensures that your `nodes.config` file is kept up-to-date with minimal manual intervention. The script dynamically adjusts the cluster configuration by reflecting changes in the security group, making the process seamless. Efficient and automated—just the way it should be!
upvoted 1 times

 **Cloudxie** 6 months, 1 week ago

D is the best

upvoted 1 times

🗲️ 👤 **Dushank** 1 year, 5 months ago

Selected Answer: A

1

The best solution to meet the company's requirements is to use AWS OpsWorks Stacks to layer the server nodes of the cluster. Create a Chef recipe that populates the content of the `/etc/cluster/nodes.config` file and restarts the service by using the current members of the layer. Assign that recipe to the Configure lifecycle event.

upvoted 2 times

🗲️ 👤 **rhinozD** 1 year, 8 months ago

Selected Answer: A

A is correct.

This event occurs on all of the stack's instances when one of the following occurs:

An instance enters or leaves the online state.

You associate an Elastic IP address with an instance or disassociate one from an instance.

You attach an Elastic Load Balancing load balancer to a layer, or detach one from a layer.

<https://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-events.html>

upvoted 2 times

🗲️ 👤 **devnv** 1 year, 10 months ago

A is correct

upvoted 2 times

A DevOps engineer is working on a data archival project that requires the migration of on-premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location. The script uses the S3 PutObject operation.

During a code review, the DevOps engineer notices that the script does not verify whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.

Which solutions for the script will meet these requirements? (Choose two.)

- A. Check the returned response for the VersionId. Compare the returned VersionId against the MD5 checksum.
- B. Include the MD5 checksum within the Content-MD5 parameter. Check the operation call's return status to find out if an error was returned.
- C. Include the checksum digest within the tagging parameter as a URL query parameter.
- D. Check the returned response for the ETag. Compare the returned ETag against the MD5 checksum.
- E. Include the checksum digest within the Metadata parameter as a name-value pair. After upload, use the S3 HeadObject operation to retrieve metadata from the object.

Suggested Answer: BD

Community vote distribution

BD (100%)

 **haazybanj** Highly Voted 1 year, 1 month ago

Selected Answer: BD

B. Explanation: When using the S3 PutObject operation, you can include the MD5 checksum of the object in the Content-MD5 parameter of the request. Amazon S3 will calculate the MD5 checksum of the object and compare it to the provided checksum. If the checksums do not match, Amazon S3 will return an error response, indicating that the data integrity check failed. This way, you can ensure that the data was successfully copied to Amazon S3 without corruption.

D. Explanation: When you use the S3 PutObject operation, it returns an ETag in the response, which is the MD5 checksum of the object that was stored in Amazon S3. After performing the upload, you can check the returned ETag against the MD5 checksum you have locally calculated. If they match, it means the data was transferred successfully without corruption. If they don't match, it indicates a data integrity issue, and you can take appropriate actions.

upvoted 11 times

 **dzn** Most Recent 6 months, 2 weeks ago

Selected Answer: BD

If the object was created by a PutObject, PostObject, or Copy operation, or via the AWS Management Console, and the object is either plain text or encrypted with server-side encryption using the Amazon S3 managed key (SSE-S3), the object's ETag is the MD5 digest of the object data.

upvoted 3 times

 **thanhnv142** 7 months, 1 week ago

Selected Answer: BD

B and D are correct: <verify whether the data was successfully copied to Amazon S3> means we need to check <operation call's return status> code. <use MD5 checksums to verify data integrity> means we need to check ETag

A: no mention of ETag

C and E: no mention of ETag or return status code

upvoted 3 times



 **rhinozD** 1 year, 2 months ago

Selected Answer: BD

BD

refer this link: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/checking-object-integrity.html>

upvoted 4 times

  **devnv** 1 year, 3 months ago

BD are correct

upvoted 3 times

A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process, the company exports the JavaScript SDK for the API from the API Gateway console and uploads the SDK to an Amazon S3 bucket.

The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin. Web clients then download the SDK by using the CloudFront distribution's endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.

Which solution will meet these requirements?

- A. Create a CodePipeline action immediately after the deployment stage of the API. Configure the action to invoke an AWS Lambda function. Configure the Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket, and create a CloudFront invalidation for the SDK path.
- B. Create a CodePipeline action immediately after the deployment stage of the API. Configure the action to use the CodePipeline integration with API Gateway to export the SDK to Amazon S3. Create another action that uses the CodePipeline integration with Amazon S3 to invalidate the cache for the SDK path.
- C. Create an Amazon EventBridge rule that reacts to UpdateStage events from aws.apigateway. Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket, and call the CloudFront API to create an invalidation for the SDK path.
- D. Create an Amazon EventBridge rule that reacts to CreateDeployment events from aws.apigateway. Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket, and call the S3 API to invalidate the cache for the SDK path.

Suggested Answer: A

Community vote distribution

A (100%)

 **thanhnv142** Highly Voted 7 months, 1 week ago

Selected Answer: A

A is correct: <by using an AWS CodePipeline pipeline> means we need CodePipeline.

C and D: no mention of CodePipeline.

B: < Configure the action to use the CodePipeline integration with API Gateway to export the SDK to Amazon S3>: codepipeline need to invoke other tools to do its task. There is not integration with API gateway

upvoted 5 times

 **zanhsieh** Most Recent 10 months, 2 weeks ago

Selected Answer: A

Vote A. Reasons:

C: No. "aws.apigateway needs API Gateway AWS integration to send events to EventBridge without using compute service, such as Lambda or Amazon EC2."

<https://aws.amazon.com/blogs/compute/capturing-client-events-using-amazon-api-gateway-and-amazon-eventbridge/>

B&D: No. S3 API doesn't contain invalidate cache call, whereas CloudFront does. Search "invalidat" in

<https://docs.aws.amazon.com/cli/latest/reference/s3api/>

<https://docs.aws.amazon.com/cli/latest/reference/cloudfront/>

upvoted 3 times

 **RVivek** 12 months ago

Selected Answer: A

B & D are wrong as it suggests to invalidate the cache fro S3


D uses event bridge rule to invoke lambada however we know Codepipeline is available and that can be used to perform the action

upvoted 2 times

 **FEEREWMWKA** 1 year ago

I pick C due to it stating during the deployment rather than at the end.



upvoted 1 times

  **ProfXsamson** 1 year, 2 months ago

Selected Answer: A



Lambda is king

upvoted 4 times

  **habros** 1 year, 2 months ago

?????

upvoted 4 times

  **devnv** 1 year, 3 months ago

A is the right answer

upvoted 3 times

A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.

A DevOps engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation, the DevOps engineer believes the failures are due to database changes not having fully propagated before the Lambda function is invoked.

How should the DevOps engineer overcome this?

- A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function.
- B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond.
- C. Add a BeforeInstall hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function.
- D. Add a ValidateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services, such as the database, are not yet ready.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **Gomer** 8 months, 4 weeks ago

Selected Answer: A

In my research, there are only TWO CodeDeploy AppSpec ifecycle event hooks for Lambda deployment:
BeforeAllowTraffic # Use to run tasks before traffic is shifted to the deployed Lambda function version.
AfterAllowTraffic # Use to run tasks after all traffic is shifted to the deployed Lambda function version.

A: (YES) Don't redirect traffic untill ready
B: (NO) Block traffic until ready
C: (NO) Event hook for Lambda
D: (NO) Event hook for Lambda
upvoted 4 times

🗳️ **seetpt** 10 months, 1 week ago

Selected Answer: A

I think A
upvoted 1 times

🗳️ **jojom19980** 1 year ago

Selected Answer: A

D can be correct if there is a wait to database to be ready so I will go with A
upvoted 3 times

🗳️ **thanhv142** 1 year, 1 month ago



A is correct: <using AWS CodeDeploy to deploy> and <a CI/CD pipeline> means lifecycle event hook
B: AfterAllowTraffic wont solve the problem, we need to hook before traffic is allowed, as in <not having fully propagated before the Lambda function is invoked>
C: beforeInstall is used to prepare for the installation process, so it is not relevant
D: there is no ValidateService hook
upvoted 4 times

🗳️ **habros** 1 year, 8 months ago

"Hooks": [
{

```
"BeforeInstall": "BeforeInstallHookFunctionName"
},
{
  "AfterInstall": "AfterInstallHookFunctionName"
},
{
  "AfterAllowTestTraffic": "AfterAllowTestTrafficHookFunctionName"
},
{
  "BeforeAllowTraffic": "BeforeAllowTrafficHookFunctionName"
},
{
  "AfterAllowTraffic": "AfterAllowTrafficHookFunctionName"
}
]
}
```

upvoted 3 times

  **habros** 1 year, 8 months ago

Opting for A based on this

upvoted 2 times

  **HugoFM** 1 year, 3 months ago

Those hooks are not valid for a Lambda, see the doc.

Lambda only supports BeforeAllowTraffic and AfterAllowTraffic. Anyway the answer is A



upvoted 6 times

  **[Removed]** 1 year, 8 months ago

Selected Answer: A



A is make sense

upvoted 3 times

  **OrganizedChaos25** 1 year, 9 months ago

A is the right answer

upvoted 2 times

  **devnv** 1 year, 10 months ago

A is correct

upvoted 2 times

A company uses a single AWS account to test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted-ssh AWS Config managed rule.

The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.

A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.

What should the DevOps engineer do next to meet these requirements?

- A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge rule. Configure the EventBridge rule to publish a notification to the SNS topic.
- B. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topic. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.
- C. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic.
- D. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT. Configure an input transformer for the restricted-ssh rule. Configure the EventBridge rule to publish a notification to the SNS topic.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **steli0** 3 months, 2 weeks ago

Selected Answer: A

D is tricky since it's not clear if the input transformer mentioned in the answer is supposed to be applied to the config rule or the EventBridge rule.

upvoted 1 times

🗳️ **zijo** 9 months, 1 week ago

AWS Config can send notifications to an SNS topic directly but here you need a customized notification which is only possible with the input transformer in Amazon EventBridge. So I think A is the better choice.

upvoted 3 times

🗳️ **MalonJay** 10 months, 1 week ago

B

AWS Config can send notifications directly to SNS.

upvoted 2 times

🗳️ **Heyang** 1 year ago

🔗 About strict-ssh https://docs.aws.amazon.com/zh_cn/config/latest/developerguide/restricted-ssh.html

upvoted 1 times

🗳️ **thanhnv142** 1 year, 1 month ago

Selected Answer: A

A is correct: <needs an automated monitoring solution that will provide a customized notification> and <creates an Amazon Simple Notification Service (Amazon SNS) topic> means they have already have SNS. we need to trigger alarm with eventbridge and send noti to SNS

B: no mention of event bride

C: AWS Systems Manager Run Command on the SNS topic to customize a notification: this step is unnecessary

D: <matches all AWS Config evaluation results of NON_COMPLIAN>: we need to match NON_COMPLIANT for the restricted-ssh rule only

upvoted 4 times



🗳️ **beanxyz** 1 year, 6 months ago

Selected Answer: A

Here is an example



<https://repost.aws/knowledge-center/config-resource-non-compliant>

upvoted 2 times

  **Aja1** 1 year, 7 months ago

Option C is the most appropriate solution for creating a customized SNS notification when the restricted-ssh AWS Config rule is evaluated as NON_COMPLIANT.

upvoted 1 times

  **Aja1** 1 year, 6 months ago

Sorry A

EventBridge input transformers are used to customize the data that is sent to a target of an EventBridge rule. They can be used to extract specific data from the event, to convert the data to a different format, or to filter the data.

upvoted 3 times

  **Jaguaroooo** 1 year, 2 months ago

why would you want to customize anything to SNS. I chose C, but A makes more sense. no need for sns customization

upvoted 1 times

  **haazybanj** 1 year, 7 months ago

Selected Answer: A


A

The Amazon EventBridge rule should be set up to match AWS Config evaluation results specifically for the restricted-ssh rule.

An input transformer should be configured for the EventBridge rule to extract and format the required information (e.g., name and ID of the noncompliant security group) from the AWS Config evaluation result.

The EventBridge rule should be configured to publish a notification to the SNS topic once it detects a noncompliant result.

upvoted 3 times

  **[Removed]** 1 year, 8 months ago



Selected Answer: A

A is good,

- restrict trigger to only ssh sg non compliance

- you need input trans*** for sending message to SNS

upvoted 4 times

  **devnv** 1 year, 10 months ago

A is the right answer

upvoted 2 times

A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.

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

- A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store. Use Aurora's automatic recovery capabilities in the event of a disaster.
- B. Create an Amazon Aurora global database in two Regions as the data store. In the event of a failure, promote the secondary Region as the primary for the application.
- C. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store. Use a Network Load Balancer to balance the database traffic in different Regions.
- D. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions. Use health checks to determine the availability in a given Region. Use Auto Scaling groups in each Region to adjust capacity based on demand.
- E. Set up the application in two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand. In the event of a disaster, adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Suggested Answer: BD

Community vote distribution

BD (100%)

🗳️ 👤 **zijo** 3 months ago

Amazon Aurora clusters are designed to be region-specific, meaning that an Aurora DB cluster is limited to a single AWS region. However, you can use Amazon Aurora Global Database to span multiple AWS regions. But Amazon Aurora clusters can span across multiple AZs.

An AWS Auto Scaling group cannot span multiple regions. Each Auto Scaling group is limited to a single AWS region. However, within that region, an Auto Scaling group can span multiple Availability Zones to ensure high availability and fault tolerance.

upvoted 1 times

🗳️ 👤 **thanhnv142** 7 months, 1 week ago

Selected Answer: BD

B and D are correct: <needs a strategy for failover and disaster recovery> means global table or db cluster and route53 fail-over policy

A and C: These options mention spanning an Amazon Aurora cluster across multiple region. This is not true. A cluster can span across multiple AZs, not regions. The only Aurora solution that can span multiple regions is global table, which includes multiple clusters.

E: No mention of Route 53 failover-based routing

upvoted 4 times

🗳️ 👤 **MaiHuong** 11 months ago

between ABC, choose B. A is wrong because "Amazon Aurora cluster in one Availability Zone across multiple Regions" is nonsense. C is incorrect too because Aurora multi-master cluster can't be across multiple regions

between DE, choose D because using Route 53 failover-based routing makes sense. E is wrong Auto Scaling group can't be multi-region

upvoted 4 times

🗳️ 👤 **Snape** 1 year, 1 month ago

Selected Answer: BD

No brainer

upvoted 4 times

🗳️ 👤 **OrganizedChaos25** 1 year, 3 months ago

Got BD as my answers

upvoted 4 times

🗳️ 👤 **devnv** 1 year, 3 months ago

BD are correct

upvoted 2 times

  **ParagSanyashiv** 1 year, 4 months ago

Selected Answer: BD

BD is the correct answer

upvoted 3 times

A business has an application that consists of five independent AWS Lambda functions.

The DevOps engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds, tests, packages, and deploys each Lambda function in sequence. The pipeline uses an Amazon EventBridge rule to ensure the pipeline starts as quickly as possible after a change is made to the application source code.

After working with the pipeline for a few months, the DevOps engineer has noticed the pipeline takes too long to complete.

What should the DevOps engineer implement to BEST improve the speed of the pipeline?

- A. Modify the CodeBuild projects within the pipeline to use a compute type with more available network throughput.
- B. Create a custom CodeBuild execution environment that includes a symmetric multiprocessing configuration to run the builds in parallel.
- C. Modify the CodePipeline configuration to run actions for each Lambda function in parallel by specifying the same runOrder.
- D. Modify each CodeBuild project to run within a VPC and use dedicated instances to increase throughput.

Suggested Answer: C

Community vote distribution

C (100%)

  **Dushank**  11 months, 3 weeks ago

Selected Answer: C

Parallel Execution:

By modifying the CodePipeline configuration to run actions for each Lambda function in parallel with the same runOrder, you allow multiple Lambda functions to be built and deployed simultaneously, which significantly improves the overall speed of the pipeline.

RunOrder:




The runOrder parameter in CodePipeline allows you to specify the order in which actions run. If multiple actions have the same runOrder, they can run in parallel.

upvoted 9 times

  **davdan99** 8 months ago

Thanks for runOrder




upvoted 1 times

  **ParagSanyashiv**  1 year, 4 months ago

Selected Answer: C

Agree with C

upvoted 7 times



  **thanhnv142**  7 months, 1 week ago

C is correct: <the pipeline takes too long to complete> and <consists of five independent AWS Lambda functions> means we should run the lambda funcs in parallel by specifying the same runOrder

A and D: no mention of running in parallel

B: No mention of runOrder

upvoted 2 times

  **MarDog** 1 year, 2 months ago

Selected Answer: C



Yeah, it's definitely C.

upvoted 4 times

  **OrganizedChaos25** 1 year, 3 months ago

Answer is C

upvoted 3 times

  **devnv** 1 year, 3 months ago

C is right answer
upvoted 3 times

A company uses AWS CloudFormation stacks to deploy updates to its application. The stacks consist of different resources. The resources include AWS Auto Scaling groups, Amazon EC2 instances, Application Load Balancers (ALBs), and other resources that are necessary to launch and maintain independent stacks. Changes to application resources outside of CloudFormation stack updates are not allowed.

The company recently attempted to update the application stack by using the AWS CLI. The stack failed to update and produced the following error message: "ERROR: both the deployment and the CloudFormation stack rollback failed. The deployment failed because the following resource(s) failed to update: [AutoScalingGroup]."

The stack remains in a status of UPDATE_ROLLBACK_FAILED.

Which solution will resolve this issue?

- A. Update the subnet mappings that are configured for the ALBs. Run the `aws cloudformation update-stack-set` AWS CLI command.
- B. Update the IAM role by providing the necessary permissions to update the stack. Run the `aws cloudformation continue-update-rollback` AWS CLI command.
- C. Submit a request for a quota increase for the number of EC2 instances for the account. Run the `aws cloudformation cancel-update-stack` AWS CLI command.
- D. Delete the Auto Scaling group resource. Run the `aws cloudformation rollback-stack` AWS CLI command.

Suggested Answer: B

Community vote distribution

B (100%)

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


Selected Answer: B

<https://repost.aws/knowledge-center/cloudformation-update-rollback-failed>

If your stack is stuck in the UPDATE_ROLLBACK_FAILED state after a failed update, then the only actions that you can perform on the stack are the ContinueUpdateRollback or DeleteStack operations.


So only B has ContinueUpdateRollback

upvoted 11 times

 **zijo** Most Recent 3 months ago

To update an AWS CloudFormation stack, you need an IAM role with permissions that allow you to perform the necessary actions on the resources defined in your CloudFormation template, as well as on the CloudFormation service itself. B is the answer


upvoted 1 times

 **thanhnv142** 7 months, 1 week ago

B is correct: <UPDATE_ROLLBACK_FAILED> means we are left with only two options: continue-update-rollback or delete-stack. We should provide necessary permissions to update the stack as well

A, C and D: no mention of continue-update-rollback or adding necessary permissions


upvoted 3 times

 **Dushank** 11 months, 3 weeks ago

Selected Answer: B

They should update the IAM role by providing the necessary permissions to update the stack and then run the `aws cloudformation continue-update-rollback` AWS CLI command

upvoted 2 times

 **Blueee** 1 year, 2 months ago

Selected Answer: B



B is correct

upvoted 3 times

 **OrganizedChaos25** 1 year, 3 months ago

B is the answer I got

upvoted 3 times

  **devnv** 1 year, 3 months ago

B is correct

upvoted 3 times

A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity.



Which solution will meet these requirements?

- A. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs. Configure AWS CloudTrail to deliver the API logs to Amazon S3. Use CloudWatch to query both sets of logs.
- B. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs. Configure AWS CloudTrail to deliver the API logs to CloudWatch Logs. Use CloudWatch Logs Insights to query both sets of logs.
- C. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon Kinesis. Configure AWS CloudTrail to deliver the API logs to Kinesis. Use Kinesis to load the data into Amazon Redshift. Use Amazon Redshift to query both sets of logs.
- D. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon S3. Use AWS CloudTrail to deliver the API logs to Amazon S3. Use Amazon Athena to query both sets of logs in Amazon S3.

Suggested Answer: B

Community vote distribution

B (100%)

  **ogwu2000** Highly Voted 1 year, 1 month ago



B is correct

A - wrong because CloudWatch is not a query tool.

C - Wrong because CloudWatch agent cant send logs directly to Kinesis. Should be from CloudWatch log

D - Wrong because CloudWatch agent cant send logs directly to S3. Should be from CloudWatch log to firehose to S3

upvoted 7 times

  **thanhnv142** Most Recent 7 months, 1 week ago

B is correct: <query application logs and AWS account API activity> means we need cloudwatch log and cloud trail

C and D: cloudwatch agent cannot directly send logs to S3 or Kinesis.

A: Cloudwatch query works only on cloudwatch, not S3

upvoted 4 times

  **haazybanj** 1 year, 1 month ago

Selected Answer: B

Explanation:

Option B provides a comprehensive solution for querying application logs and AWS account API activity. The Amazon CloudWatch agent is used to send logs from the EC2 instances to Amazon CloudWatch Logs, allowing easy access to application logs. AWS CloudTrail is configured to deliver the API logs to CloudWatch Logs, enabling monitoring and analysis of AWS account activity. Finally, CloudWatch Logs Insights is utilized to query and analyze both sets of logs efficiently.

upvoted 4 times

  **ProfXsamson** 1 year, 2 months ago

Selected Answer: B



Since Cloudwatch Insights can perform query, no need to use s3/athena.

upvoted 4 times

  **OrganizedChaos25** 1 year, 3 months ago

B is correct

upvoted 1 times

  **devnv** 1 year, 3 months ago

B is the right answer

upvoted 1 times

A company wants to ensure that their EC2 instances are secure. They want to be notified if any new vulnerabilities are discovered on their instances, and they also want an audit trail of all login activities on the instances.

Which solution will meet these requirements?

- A. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances. Install the Amazon Kinesis Agent to capture system logs and deliver them to Amazon S3.
- B. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances. Install the Systems Manager Agent to capture system logs and view login activity in the CloudTrail console.
- C. Configure Amazon CloudWatch to detect vulnerabilities on the EC2 instances. Install the AWS Config daemon to capture system logs and view them in the AWS Config console.
- D. Configure Amazon Inspector to detect vulnerabilities on the EC2 instances. Install the Amazon CloudWatch Agent to capture system logs and record them via Amazon CloudWatch Logs.

Suggested Answer: D

Community vote distribution

D (100%)

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

Selected Answer: D

Aamzon Inspector detect software vulnerabilities and unintended network exposure in near real time in AWS workloads such as Amazon EC2, AWS Lambda functions, and Amazon ECR.

upvoted 9 times

🗳️ 👤 **thanhnv142** Highly Voted 1 year, 1 month ago

Selected Answer: D

D is correct: <new vulnerabilities are discovered> means AWS inspector

A and B: AWS SSM does not support vulnerabilities scanning

C: Amazon CloudWatch does not support vulnerabilities scanning

upvoted 5 times

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

Selected Answer: D

Exam tip : if you see "vulnerabilities" in the question -> choose the answer covers "Amazon Inspector"

upvoted 3 times

🗳️ 👤 **EVAAWS** 7 months ago

Selected Answer: D

D is the right answer

upvoted 1 times

🗳️ 👤 **OrganizedChaos25** 1 year, 9 months ago

D is correct

upvoted 2 times

🗳️ 👤 **devnv** 1 year, 10 months ago

D is the right answer

upvoted 2 times

A company is running an application on Amazon EC2 instances in an Auto Scaling group. Recently, an issue occurred that prevented EC2 instances from launching successfully, and it took several hours for the support team to discover the issue. The support team wants to be notified by email whenever an EC2 instance does not start successfully.

Which action will accomplish this?

- A. Add a health check to the Auto Scaling group to invoke an AWS Lambda function whenever an instance status is impaired.
- B. Configure the Auto Scaling group to send a notification to an Amazon SNS topic whenever a failed instance launch occurs.
- C. Create an Amazon CloudWatch alarm that invokes an AWS Lambda function when a failed AttachInstances Auto Scaling API call is made.
- D. Create a status check alarm on Amazon EC2 to send a notification to an Amazon SNS topic whenever a status check fail occurs.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **MarDog** Highly Voted 👍 1 year, 2 months ago

Selected Answer: B

Likely B:

<https://aws.amazon.com/blogs/aws/auto-scaling-notifications-recurrence-and-more-control/>

"EC2_INSTANCE_LAUNCH_ERROR"

upvoted 7 times

🗳️ 👤 **thanhnv142** Most Recent 🕒 7 months, 1 week ago

B is correct: <wants to be notified by email> means SNS. <EC2 instances in an Auto Scaling group> means this should be triggered by auto scaling group

A, C and D: no mention of both SNS and auto scaling group

upvoted 4 times

🗳️ 👤 **rhinozD** 1 year, 2 months ago

Selected Answer: B

B is correct

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

upvoted 3 times

🗳️ 👤 **OrganizedChaos25** 1 year, 3 months ago

I got B as my answer

upvoted 1 times

🗳️ 👤 **2pk** 1 year, 3 months ago

Selected Answer: B

i think B is correct , but Option A is incorrect because, this would only be triggered if an instance was already running and experiencing issues. It would not provide notification when an instance fails to launch.

upvoted 3 times

🗳️ 👤 **devnv** 1 year, 3 months ago

B is correct

upvoted 1 times

A company is using AWS Organizations to centrally manage its AWS accounts. The company has turned on AWS Config in each member account by using AWS CloudFormation StackSets. The company has configured trusted access in Organizations for AWS Config and has configured a member account as a delegated administrator account for AWS Config.


A DevOps engineer needs to implement a new security policy. The policy must require all current and future AWS member accounts to use a common baseline of AWS Config rules that contain remediation actions that are managed from a central account. Non-administrator users who can access member accounts must not be able to modify this common baseline of AWS Config rules that are deployed into each member account.

Which solution will meet these requirements?

- A. Create a CloudFormation template that contains the AWS Config rules and remediation actions. Deploy the template from the Organizations management account by using CloudFormation StackSets.
- B. Create an AWS Config conformance pack that contains the AWS Config rules and remediation actions. Deploy the pack from the Organizations management account by using CloudFormation StackSets.
- C. Create a CloudFormation template that contains the AWS Config rules and remediation actions. Deploy the template from the delegated administrator account by using AWS Config.
- D. Create an AWS Config conformance pack that contains the AWS Config rules and remediation actions. Deploy the pack from the delegated administrator account by using AWS Config.

Suggested Answer: D

Community vote distribution



Certified101 Highly Voted 1 year, 2 months ago

Selected Answer: D

Option D. Create an AWS Config conformance pack that contains the AWS Config rules and remediation actions. Deploy the pack from the delegated administrator account by using AWS Config.

Conformance packs are a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a region, and across an organization in AWS Organizations. These packs are created and managed from a central account, and help to establish a secure and compliant posture for your accounts. Non-administrator users can view the AWS Config rules within a conformance pack but they cannot modify them. AWS Config conformance packs are therefore a good fit for achieving the desired control and security policy.

The other options, while potentially viable for deploying Config rules, do not inherently protect the baseline AWS Config rules from being modified by non-administrator users in the member accounts.

upvoted 10 times

thanhv142 Highly Voted 7 months, 1 week ago

Selected Answer: D

D is correct: <a common baseline of AWS Config rule> means conformance pack. <a member account as a delegated administrator account for AWS Config> means delegated admin

A and C: no mention of conformance pack

B: should deploy this using AWS config and in the delegated account, not the management account

upvoted 5 times

MalonJay Most Recent 4 months ago

The question says

'The policy must require all current and future AWS member accounts to use a common baseline of AWS Config rules'

Does D account for that?

upvoted 1 times

lunt 1 year, 2 months ago

Selected Answer: D

Not sure why some people are saying B.

A= CFN cannot protect the config.

B= Yes technically, where is the actual CONFIG management plane? Its in the delegated admin account, which is not the management account = delegated admin config account will have no idea of management account config.

C= CFN cannot protect config.

D= Yes. Delegated CONFIG account can config on orgz level & protect the rules. Only logical option.

upvoted 5 times

🗳️ 👤 **allen_devops** 1 year, 2 months ago

D is correct. Deploying via Cloudformation StackSet cannot make sure that the aws config itself is not modified by the member accounts. Deploy aws organizational rule will achieve both permission restriction and auto deployment

<https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html>

upvoted 1 times

🗳️ 👤 **rhinozD** 1 year, 2 months ago

Selected Answer: D

D is correct

<https://aws.amazon.com/blogs/mt/deploying-conformance-packs-across-an-organization-with-automatic-remediation/>

upvoted 3 times

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

option B is the most appropriate solution for centrally managing and enforcing the common baseline of AWS Config rules across all member accounts while ensuring that non-administrator users cannot modify the rules.

upvoted 1 times

🗳️ 👤 **stream3652** 1 year, 3 months ago

Can't you use D?

upvoted 2 times

🗳️ 👤 **2pk** 1 year, 3 months ago

Selected Answer: B

i think its B, because AWS Config conformance packs are a way to package AWS Config rules and remediation actions into a single, shareable entity. With AWS Organizations, you can use CloudFormation StackSets to deploy conformance packs across all member accounts in your organization. This allows you to centrally manage the deployment of AWS Config rules and remediation actions across multiple AWS accounts. By deploying the conformance pack from the Organizations management account, you can ensure that non-administrator users cannot modify the baseline rules deployed to each member account.

upvoted 3 times

🗳️ 👤 **Jaguaroooo** 8 months, 1 week ago

the question tells you there's a "delegated account". so your answer should be looking for that account in your answer choices as well.

upvoted 2 times

🗳️ 👤 **rhinozD** 1 year, 2 months ago

No, you just need the manager account to deploy the conformance pack to all organization.

upvoted 1 times

🗳️ 👤 **emupsx1** 1 year, 2 months ago

It's B

<https://catalog.us-east-1.prod.workshops.aws/workshops/7bb9fd8f-d049-4163-98e3-5c0cbb211f0c/en-US/enable-custom-conformance-pack-using-stacksets>

upvoted 1 times

🗳️ 👤 **devnv** 1 year, 3 months ago

D is the right answer

upvoted 3 times

🗳️ 👤 **Jeanphi72** 1 year, 4 months ago

Selected Answer: D

<https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html>

upvoted 5 times

A DevOps engineer manages a large commercial website that runs on Amazon EC2. The website uses Amazon Kinesis Data Streams to collect and process web logs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2.


Sudden increases of data cause the Kinesis consumer application to fall behind, and the Kinesis data streams drop records before the records can be processed. The DevOps engineer must implement a solution to improve stream handling.

Which solution meets these requirements with the MOST operational efficiency?

- A. Modify the Kinesis consumer application to store the logs durably in Amazon S3. Use Amazon EMR to process the data directly on Amazon S3 to derive customer insights. Store the results in Amazon S3.
- B. Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch GetRecords.IteratorAgeMilliseconds metric. Increase the retention period of the Kinesis data streams.
- C. Convert the Kinesis consumer application to run as an AWS Lambda function. Configure the Kinesis data streams as the event source for the Lambda function to process the data streams.
- D. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Suggested Answer: B

Community vote distribution



🗳️ 👤 **emupsx1** Highly Voted 1 year, 7 months ago

The answer is B because:

A few hours ago, I just finished the DOP-C02 exam.

My score is 1000 points.

This question has come up, I choose B.

upvoted 41 times

🗳️ 👤 **youonebe** 2 months, 2 weeks ago

Remember, only 50 of the questions are graded.

First, there is no way what you said is true.

Second, your answer might not be graded.

Correct answer is a firm C.

upvoted 3 times

🗳️ 👤 **[Removed]** 1 year, 5 months ago

Were there any other questions from here in the exam?

upvoted 2 times

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

First of all Congratulations.

Now how do you know that this question was not from those 10 questions that do not count towards your score. and your answer to this question was Wrong but not counted towards your score. Just Saying!!

Peace :)

upvoted 1 times

🗳️ 👤 **Jaguaroooo** 1 year, 2 months ago

B is the Answer, let him show off, it is ok.

upvoted 5 times

🗳️ 👤 **yorkicurke** Highly Voted 1 year, 3 months ago

Selected Answer: B

why not C?

because we just replace ONE ec2 with ONE lambda here. And no mention of aws lambda reserved concurrency or provisioned concurrency.

In the question were are asked for 'MOST operational efficiency'. that's my two cents.

Ciao

upvoted 6 times



  **ce0df07** Most Recent 1 month ago

Selected Answer: B

Option B because:

1. It uses `GetRecords.IteratorAgeMilliseconds` metric
 - This metric indicates how far behind the consumer is processing
 - Perfect indicator for when scaling is needed
 - Automatically detects processing lag
2. Horizontal scaling of EC2 consumers
 - Adds processing capacity when needed
 - Can scale down when demand decreases
 - Maintains existing application architecture
3. Increasing retention period
 - Provides buffer against temporary processing delays
 - Prevents data loss during scaling events
 - Default is 24 hours, can be increased up to 365 days

upvoted 1 times

  **2d943d1** 1 month, 2 weeks ago

Selected Answer: C

I think most people are missing the part where the question focuses on MOST operational efficiency. Lambda fits this purpose, as we are not managing instances and Lambda can scale to meet throughput demands.

upvoted 1 times

  **ZinggieG87** 2 months, 2 weeks ago



Selected Answer: D

A, sounds a lot effort towards a different architecture solution.

B, I don't think the insufficient Kinesis throughput has can be resolved.

C, adding lambda is a big change, no mention about the concurrency limit isn't clear mentioned in the question.

upvoted 1 times

  **youonebe** 2 months, 2 weeks ago

Selected Answer: C

Answer is C.

B - This solution improves the consumer application's ability to handle increased throughput, but it still requires manual scaling of EC2 instances, which can involve some complexity and operational overhead.

C - Using Lambda is obviously more Operational Efficient.

upvoted 1 times

  **Gomer** 8 months, 3 weeks ago

Selected Answer: B

I think B makes more sense, because "the Kinesis data streams drop records before the records can be processed". It's not an intake throughput issue that needs more shards, but an outtake throughput issue.

"Consumer Record Processing Falling Behind"

"After you identify how far behind your consumers are reading, look at the most common reasons why consumers fall behind. Start with the `GetRecords.IteratorAgeMilliseconds` metric, which tracks the read position across all shards and consumers in the stream. Note that if an iterator's age passes 50% of the retention period (by default, 24 hours...), there is risk for data loss due to record expiration."

"A quick stopgap solution is to increase the retention period."

[...]

"An alternative approach is to increase your parallelism by increasing the number of shards."

"Finally, confirm you have an adequate amount of physical resources (memory, CPU utilization, etc.) on the underlying processing nodes during peak demand."

<https://docs.aws.amazon.com/streams/latest/dev/troubleshooting-consumers.html>

upvoted 2 times

  **zijo** 9 months, 1 week ago

B is a good choice

The `GetRecords.IteratorAgeMilliseconds` metric in Amazon CloudWatch for Amazon Kinesis Data Streams measures the age of the last record returned by the `GetRecords` operation. Specifically, it represents the time difference between the current time and the approximate arrival timestamp of the last record processed by a consumer in milliseconds.

Purpose: Measures the delay in processing data records from the time they are added to the stream to the time they are processed by a consumer.

Scaling: If you observe high iterator age values, consider increasing the number of shards or enhancing the processing capacity of your consumers (e.g., adding more instances or increasing the processing power of existing instances).

upvoted 1 times

🗳️ 👤 **seetpt** 10 months, 1 week ago

Selected Answer: B

I choose B

upvoted 1 times

🗳️ 👤 **c3518fc** 10 months, 3 weeks ago

Selected Answer: B

Consumer Record Processing Falling Behind

For most use cases, consumer applications are reading the latest data from the stream. In certain circumstances, consumer reads may fall behind, which may not be desired. After you identify how far behind your consumers are reading, look at the most common reasons why consumers fall behind.

Start with the `GetRecords.IteratorAgeMilliseconds` metric, which tracks the read position across all shards and consumers in the stream. Note that if an iterator's age passes 50% of the retention period (by default, 24 hours, configurable up to 365 days), there is risk for data loss due to record expiration. A quick stopgap solution is to increase the retention period. This stops the loss of important data while you troubleshoot the issue further. <https://docs.aws.amazon.com/streams/latest/dev/troubleshooting-consumers.html#record-processing-falls-behind>

upvoted 4 times

🗳️ 👤 **Shasha1** 1 year ago

B

`GetRecords.IteratorAgeMilliseconds` metric : its for track the progress of Kineses consumer, this cloud watch matric is use for the measuring the difference between current time and when the last record of `GetRecords` calls written to the stream. `IteratorAgeMilliseconds` metric is 0 means is processing fast enough and if its >0 its slow in processing. Therefore, Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch `GetRecords.IteratorAgeMilliseconds` metric. Increase the retention period of the Kinesis data streams.

upvoted 2 times

🗳️ 👤 **kyuhuck** 1 year ago

Selected Answer: C

option C (Converting the Kinesis consumer application to run as an AWS Lambda function) is the most suitable solution. This approach automatically scales with the amount of incoming data, reduces the operational burden of managing EC2 instances, and leverages the serverless model to only incur costs for the actual compute time used for processing the data. This option provides a scalable, efficient, and cost-effective solution to the problem without the need for extensive infrastructure management.

upvoted 2 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

B is correct: <consumer application to fall behind> means we need to increase the power of the consumer. <Kinesis data streams drop records> means we should extend timeout. Only B match these requirements

A: irrelevant

C: Lambda is used only for short-lived tasks because its maximum execution time is 15 min. In this case, we need to process web logs. This is a time-consuming task, which is not suitable for Lambda

D: No mention of increasing Consumer power

upvoted 1 times

🗳️ 👤 **svjl** 1 year, 3 months ago

B & D are correct. However, the question is looking for a solution for two issues

"application to fall behind, and the Kinesis data streams drop records before the records can be processed"

Then, B is the most appropriate solution



upvoted 3 times

🗳️ 👤 **2pk** 1 year, 4 months ago

B is the answer: The data fall behind due to lack of physical resources at consumer side. Increase of more nodes will address this issue.

<https://docs.aws.amazon.com/streams/latest/dev/troubleshooting-consumers.html#record-processing-falls-behind>

upvoted 1 times

  **buenos** 1 year, 5 months ago

Selected Answer: B

I would say B

upvoted 2 times

  **Dushank** 1 year, 5 months ago

Selected Answer: D

D. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Here's the rationale for choosing this option:

Increasing Shards for Throughput:

By increasing the number of shards in the Kinesis data streams, you increase the overall throughput and the capacity to handle sudden increases in data. This directly addresses the issue of the consumer application falling behind during data spikes.

Operational Efficiency:

Scaling the shards provides a more straightforward and efficient solution in terms of operation compared to modifying the consumer application, horizontally scaling instances, or converting the application to run as an AWS Lambda function.

upvoted 1 times

A company recently created a new AWS Control Tower landing zone in a new organization in AWS Organizations. The landing zone must be able to demonstrate compliance with the Center for Internet Security (CIS) Benchmarks for AWS Foundations.


The company's security team wants to use AWS Security Hub to view compliance across all accounts. Only the security team can be allowed to view aggregated Security Hub findings. In addition, specific users must be able to view findings from their own accounts within the organization. All accounts must be enrolled in Security Hub after the accounts are created.

Which combination of steps will meet these requirements in the MOST automated way? (Choose three.)

- A. Turn on trusted access for Security Hub in the organization's management account. Create a new security account by using AWS Control Tower. Configure the new security account as the delegated administrator account for Security Hub. In the new security account, provide Security Hub with the CIS Benchmarks for AWS Foundations standards.
- B. Turn on trusted access for Security Hub in the organization's management account. From the management account, provide Security Hub with the CIS Benchmarks for AWS Foundations standards.
- C. Create an AWS IAM Identity Center (AWS Single Sign-On) permission set that includes the required permissions. Use the CreateAccountAssignment API operation to associate the security team users with the permission set and with the delegated security account.
- D. Create an SCP that explicitly denies any user who is not on the security team from accessing Security Hub.
- E. In Security Hub, turn on automatic enablement.
- F. In the organization's management account, create an Amazon EventBridge rule that reacts to the CreateManagedAccount event. Create an AWS Lambda function that uses the Security Hub CreateMembers API operation to add new accounts to Security Hub. Configure the EventBridge rule to invoke the Lambda function.

Suggested Answer: ACE

Community vote distribution



emupsx1 Highly Voted 1 year, 7 months ago

The answer is ACE because:

A few hours ago, I just finished the DOP-C02 exam.

My score is 1000 points.

This question has come up, I choose ACE.

upvoted 17 times

BaburTurk 1 year, 6 months ago

bot account, Pics or it didn't happen,

upvoted 7 times

Gomer 8 months, 3 weeks ago

Either a bot or a bot for brains. Same useless comments made on multiple questions.

upvoted 3 times

eugene2owl Most Recent 3 months, 1 week ago

Selected Answer: ADE

I prefer "D" over "C", because no-one asks to enable SSO (which is very complex to organise and maintain)

upvoted 1 times

auxwww 7 months, 2 weeks ago

Selected Answer: ACE

A - Only security team needs access to findings org wide - hence delegated account

C - Allow security team members access to delegated account for Security hub using Identity center of control tower

E - Each new account needs security hub for it's own users to access and also for aggregation across org

upvoted 3 times

zijo 9 months ago

Automatic enablement in AWS Security Hub refers to the feature that allows AWS Security Hub to be automatically enabled for new and existing AWS accounts that are part of an organization in AWS Organizations. This feature simplifies the process of onboarding multiple accounts into Security Hub, ensuring consistent security posture and compliance across the organization.

upvoted 2 times

🗳️ 👤 **seetpt** 10 months, 1 week ago

Selected Answer: ACE

ACE is correct

upvoted 2 times

🗳️ 👤 **didek1986** 10 months, 4 weeks ago

Selected Answer: ACF

ACF

E - ensures that all new accounts are automatically enrolled in Security Hub (same as F) but it does not address the requirement for specific users to view findings from their own accounts

upvoted 2 times

🗳️ 👤 **zijo** 9 months ago

I think C will take care of this. "it does not address the requirement for specific users to view findings from their own accounts"

upvoted 1 times

🗳️ 👤 **didek1986** 10 months, 4 weeks ago

ACF

E - ensures that all new accounts are automatically enrolled in Security Hub (same as F) but it does not address the requirement for specific users to view findings from their own accounts

upvoted 1 times

🗳️ 👤 **dkp** 11 months ago

Selected Answer: ACE

ace are correct answer

upvoted 2 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

ACE are correct: <Only the security team can be allowed to view aggregated Security Hub findings> means we need a delegated admin. <All accounts must be enrolled in Security Hub after the accounts are created> and <in the MOST automated way> means we need enable automatic enablement

B: no mention of delegated admin

D: This options denied access of the security team, which is irrelevant

F: This option's result is the same as in option E, but more complicated

upvoted 2 times

🗳️ 👤 **2pk** 1 year, 4 months ago

Selected Answer: ADE

According to this article .. The Delegated account users have access in ANY account while the users under own account can view their own findings. So, there is no need to setup IAM policies for Security account users.

<https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-accounts-allowed-actions.html>

upvoted 1 times

🗳️ 👤 **YR4591** 1 year, 4 months ago

Selected Answer: ACE

A - Create delegate account for the security hub

C - Give access to users to security using permissions sets

E - Use auto enable so every new account will be monitored by security hub

upvoted 3 times

🗳️ 👤 **sb333** 1 year, 7 months ago

Selected Answer: ACE

ACE are the correct answers.

upvoted 2 times

🗳️ 👤 **habros** 1 year, 8 months ago

Selected Answer: ACE

ACE. Reason being, it is a landing zone and AWS SSO (IAM IC) is already part of the Control Tower product! Add security dept users as a SSO group and attach the security permission set to access security hub

upvoted 2 times

🗨️ 👤 **Wardove** 1 year, 8 months ago

Selected Answer: ACE

with Control Tower comes the Identity Center implementation with default Identity Center directory.

upvoted 3 times

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

Selected Answer: ADE

B is not the typical way AWS separates responsibilities in multi account (management, sec, audit)

C is related with Active Directory

E is more automated than F

upvoted 2 times

🗨️ 👤 **jnv007** 1 year, 8 months ago

Identity Center is not exclusively related to Active Directory

An SCP can only prevent access but doesnt enable any access, so D is not sufficient

ACE for me

<https://docs.aws.amazon.com/securityhub/latest/userguide/accounts-orgs-auto-enable.html>

upvoted 1 times

🗨️ 👤 **Kodoma** 1 year, 9 months ago

ACF IS MORE EFFICIENT

upvoted 4 times

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

Selected Answer: ADE

Ade for me

upvoted 1 times

A company runs applications in AWS accounts that are in an organization in AWS Organizations. The applications use Amazon EC2 instances and Amazon S3.


The company wants to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future. When the company detects one of these events, the company wants to use an existing Amazon Simple Notification Service (Amazon SNS) topic to send a notification to its operational support team for investigation and remediation.

Which solution will meet these requirements in accordance with AWS best practices?

- A. In the organization's management account, configure an AWS account as the Amazon GuardDuty administrator account. In the GuardDuty administrator account, add the company's existing AWS accounts to GuardDuty as members. In the GuardDuty administrator account, create an Amazon EventBridge rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.
- B. In the organization's management account, configure Amazon GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts. Create an AWS CloudFormation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule. Configure the rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic. Configure the CloudFormation stack set to deploy into all AWS accounts in the organization.
- C. In the organization's management account, create an AWS CloudTrail organization trail. Activate the organization trail in all AWS accounts in the organization. Create an SCP that enables VPC Flow Logs in each account in the organization. Configure AWS Security Hub for the organization. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.
- D. In the organization's management account, configure an AWS account as the AWS CloudTrail administrator account. In the CloudTrail administrator account, create a CloudTrail organization trail. Add the company's existing AWS accounts to the organization trail. Create an SCP that enables VPC Flow Logs in each account in the organization. Configure AWS Security Hub for the organization. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.

Suggested Answer: B

Community vote distribution



Just_Ninja Highly Voted 1 year, 8 months ago

Selected Answer: A

Dear Admin, Please Fix the Wrong response here!

It's A:

This solution meets all the requirements:

Detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity: Amazon GuardDuty is a threat detection service that continuously monitors for malicious or unauthorized behavior. It analyzes events from AWS CloudTrail, Amazon VPC Flow Logs, and DNS logs to detect such activities.

Send a notification to the operational support team: Creating an Amazon EventBridge rule that matches GuardDuty findings and then forwarding these to an SNS topic allows for the generation of notifications whenever suspicious activity is detected.

Cover future AWS accounts: By designating a GuardDuty administrator account in AWS Organizations, you can manage GuardDuty across all of your existing and future AWS accounts. This ensures that any new account created under the organization is automatically covered by GuardDuty.

upvoted 12 times

Mrflip Most Recent 2 months ago

Selected Answer: B

B is the right answer

upvoted 1 times

jamesf 7 months, 1 week ago

Selected Answer: A

keywords: compromised EC2 instances, suspicious network activity, and unusual API activity

= GuardDuty

upvoted 2 times

🗨️ 👤 **zijo** 9 months ago

Selected Answer: A

When you use GuardDuty with an AWS organization, the management account of that organization can designate any account within the organization as the delegated GuardDuty administrator account. For this administrator account, GuardDuty gets enabled automatically only in the designated AWS Region. This account also has the permission to enable and manage GuardDuty for all of the accounts in the organization within that Region. The administrator account can view the members of and add members to this AWS organization.

AWS GuardDuty can detect unusual API activity within existing AWS accounts in an AWS Organization. It monitors AWS CloudTrail event logs, which include records of all API calls made within your AWS environment. GuardDuty analyzes these logs to identify unusual or suspicious API activity that might indicate a potential security threat.

upvoted 2 times

🗨️ 👤 **zijo** 9 months ago

A looks like a better choice.

When you use GuardDuty with an AWS organization, the management account of that organization can designate any account within the organization as the delegated GuardDuty administrator account. For this administrator account, GuardDuty gets enabled automatically only in the designated AWS Region. This account also has the permission to enable and manage GuardDuty for all of the accounts in the organization within that Region. The administrator account can view the members of and add members to this AWS organization.

upvoted 1 times

🗨️ 👤 **dkp** 11 months ago

Selected Answer: A

answer A

upvoted 2 times

🗨️ 👤 **Mordans** 11 months ago

If GuardDuty is indeed set up at the organization level (which is supported and encouraged by AWS for simplicity and coverage), then Option A becomes a very strong choice. It provides centralized management and automatic, seamless inclusion of all organization accounts in security monitoring without requiring manual intervention for each new account.

upvoted 1 times

🗨️ 👤 **stoy123** 11 months, 2 weeks ago

Selected Answer: B

Definitely B

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

Selected Answer: A

A is correct: <detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity> means AWS GuardDuty

B: dont have to invite other accounts because all accounts are in an org in AWS org.

C and D: no mention of GuardDuty

upvoted 3 times

🗨️ 👤 **a54b16f** 1 year, 1 month ago

Selected Answer: A

invitation is used to handle users OUTSIDE the organization.

upvoted 3 times

🗨️ 👤 **davdan99** 1 year, 2 months ago

Selected Answer: A

Go For A.

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_organizations.html

upvoted 2 times

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

Selected Answer: B

Member accounts must accept invite from the designated guard duty account before its effect. I use AWS organisation at work and quite familiar with the workings. I lean towards B

upvoted 2 times

🗨️ 👤 **2pk** 1 year, 4 months ago

It true it's missing auto enabled on. but Invitation is organization is not needed as Organization get precedence with account management when you have deligated Guarddduty account. "If you have already set up a GuardDuty administrator account with associated member accounts by invitation and the member accounts are part of the same organization, their Type changes from By Invitation to Via Organizations"

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_organizations.html

upvoted 1 times

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

B is the correct answer.

A could better if not for the fact that it doesn't handle automatic enablement on new AWS account.

B handles this case with CloudFormation stacksets : <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-guardduty-master.html>

upvoted 2 times

🗨️ 👤 **lunt** 1 year, 6 months ago

A= does handle automatic enablement. If the GD delegated account is setup properly with automatic enablement check box ticked. As soon as the AWS account is created, GD auto enablement kicks into gear. B = How does CFN accept the GD invite? New AWS Account. CFN runs on new account > accept GD invite...but when was this invite sent? I have to login to AWS Console > GD > create invite vs A = no invite = directly enabled for GB in new account. https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_organizations.html

upvoted 2 times

🗨️ 👤 **lunt** 1 year, 7 months ago

Selected Answer: A

B is wrong. Newly created AWS accounts = you don't need to do this if the GD Orgz is configured properly, you can accept from delegated admin account. The point remains, you can add the accounts using option A. The misdirect here is that A does not state anything about new accounts vs B which does. Bearing in mind A + B still have to do something in GD, A is actually the better option.

A is right.

upvoted 3 times

🗨️ 👤 **jason7** 1 year, 7 months ago

Selected Answer: B

Option A is not the best choice because although it correctly configures GuardDuty as the administrator, it does not handle the automatic addition of new AWS accounts to GuardDuty and the forwarding of events to the SNS topic

upvoted 1 times

🗨️ 👤 **jason7** 1 year, 7 months ago

Option B is the most suitable solution as it combines GuardDuty, AWS CloudFormation StackSets, and Amazon EventBridge to automatically monitor all existing and future AWS accounts and send notifications to the specified SNS topic when security events are detected.

upvoted 1 times

A company's DevOps engineer is working in a multi-account environment. The company uses AWS Transit Gateway to route all outbound traffic through a network operations account. In the network operations account, all account traffic passes through a firewall appliance for inspection before the traffic goes to an internet gateway.

The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO. The security team wants to receive an alert if any CRITICAL events occur.

What should the DevOps engineer do to meet these requirements?

- A. Create an Amazon CloudWatch Synthetics canary to monitor the firewall state. If the firewall reaches a CRITICAL state or logs a CRITICAL event, use a CloudWatch alarm to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email address to the topic.
- B. Create an Amazon CloudWatch metric filter by using a search for CRITICAL events. Publish a custom metric for the finding. Use a CloudWatch alarm based on the custom metric to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email address to the topic.
- C. Enable Amazon GuardDuty in the network operations account. Configure GuardDuty to monitor flow logs. Create an Amazon EventBridge event rule that is invoked by GuardDuty events that are CRITICAL. Define an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the security team's email address to the topic.
- D. Use AWS Firewall Manager to apply consistent policies across all accounts. Create an Amazon EventBridge event rule that is invoked by Firewall Manager events that are CRITICAL. Define an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the security team's email address to the topic.

Suggested Answer: B

Community vote distribution

B (100%)

 **2pk**  1 year, 3 months ago

Selected Answer: B

The logs from the firewall appliance are already being sent to Amazon CloudWatch Logs. So , The best approach to meet the given requirements is to create an Amazon CloudWatch metric filter by using a search for CRITICAL events


upvoted 7 times

 **habros**  1 year, 2 months ago

B. As the appliance pipes to CW Logs for consolidation. Define an alarm listening to the metric and should be okay.

D is ONLY CORRECT IF YOU ARE USING AWS FIREWALL MANAGER.

upvoted 5 times

 **seetpt**  4 months, 1 week ago

Selected Answer: B

I think B

upvoted 1 times

 **dkp** 4 months, 4 weeks ago

Selected Answer: B

answer B

upvoted 2 times

 **thanhnv142** 7 months, 1 week ago

B is correct: <firewall appliance sends logs to Amazon CloudWatch Logs> means we already have the log in CW logs, only need to create alarm on these log files and send to sends

A: No need to monitor the state of the firewall



C and D: no mention of CloudWatch Logs

upvoted 4 times

 **OrganizedChaos25** 1 year, 3 months ago

B is the correct answer

upvoted 1 times

  **devnv** 1 year, 3 months ago

B is correct

upvoted 1 times


A company is divided into teams. Each team has an AWS account, and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use. AWS services must gain approval through a request and approval process.

How should a DevOps engineer configure the accounts to meet these requirements?

- A. Use AWS CloudFormation StackSets to provision IAM policies in each account to deny access to restricted AWS services. In each account, configure AWS Config rules that ensure that the policies are attached to IAM principals in the account.
- B. Use AWS Control Tower to provision the accounts into OUs within the organization. Configure AWS Control Tower to enable AWS IAM Identity Center (AWS Single Sign-On). Configure IAM Identity Center to provide administrative access. Include deny policies on user roles for restricted AWS services.
- C. Place all the accounts under a new top-level OU within the organization. Create an SCP that denies access to restricted AWS services. Attach the SCP to the OU.
- D. Create an SCP that allows access to only approved AWS services. Attach the SCP to the root OU of the organization. Remove the FullAWSAccess SCP from the root OU of the organization.

Suggested Answer: C

Community vote distribution



lunt Highly Voted 1 year, 9 months ago

Selected Answer: D

A=local account admin can change this.

B=local admin has admin permissions. Complicated.

C=implicit permit on everything else = breaks requirements.

D= As they want to approve each service, its got to be white-list based SCP setup.

Answer is D.

upvoted 21 times

dkp Highly Voted 11 months ago

Selected Answer: D

Ans D:

It is easier to allow approved services than deny all the other services, considering the vast amount of AWS services. it's easier to whitelist than blacklisting all the remaining services.

upvoted 5 times

79f3aa3 Most Recent 2 weeks, 1 day ago

Selected Answer: C

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

"AWS strongly recommends that you don't attach SCPs to the root of your organization without thoroughly testing the impact that the policy has on accounts. Instead, create an OU that you can move your accounts into one at a time, or at least in small numbers, to ensure that you don't inadvertently lock users out of key services."

upvoted 1 times

ce0df07 1 month ago

Selected Answer: C

Not Option D, Using allow-list SCP, since:

- Too restrictive
- More difficult to maintain
- Might block essential services
- Could break account functionality
- Requires constant updates

upvoted 1 times

🗄️ 👤 **teo2157** 1 month, 2 weeks ago

Selected Answer: C

Going for C as removing the FullAWSAccess SCP from the root OU requires impacts directly in the Administrative Access and restrict necessary administrative actions required for account management and operations.

upvoted 2 times

🗄️ 👤 **auxwww** 4 months, 4 weeks ago

Selected Answer: D

D is more straight forward

upvoted 2 times

🗄️ 👤 **hzaki** 6 months, 2 weeks ago

Selected Answer: D

The answer is (D). The following SCP example from the AWS DOCUMENT allows accounts to create resource shares that share prefix lists https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_ram.html

upvoted 2 times

🗄️ 👤 **[Removed]** 6 months, 3 weeks ago

Selected Answer: D

Agree with D

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html

upvoted 2 times

🗄️ 👤 **jamesf** 7 months, 1 week ago

Selected Answer: C

I prefer C than D.

As SCP more in Deny but not Allow

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

upvoted 2 times

🗄️ 👤 **auxwww** 7 months, 2 weeks ago

Selected Answer: C

SCP - only deny not allow - So answer is C

upvoted 3 times

🗄️ 👤 **zsoni** 8 months, 3 weeks ago

The question is looking to use the Allow List Strategy using SCP. So the answer that best fits is D.

upvoted 1 times

🗄️ 👤 **zijo** 9 months ago

Selected Answer: C

SCPs primary function is not grant permissions by themselves but restrict the permissions that IAM policies and other access control mechanisms can grant.

upvoted 4 times

🗄️ 👤 **seetpt** 10 months, 1 week ago

Selected Answer: D

D seems better.

upvoted 2 times

🗄️ 👤 **fdoxxx** 11 months, 3 weeks ago

Selected Answer: C

Option C:

Place all the accounts under a new top-level OU within the organization: This allows for centralized management of the accounts.

Create an SCP that denies access to restricted AWS services: This ensures that only approved services are accessible. SCPs (Service Control Policies) are the best way to control permissions at the organizational level.

Attach the SCP to the OU: By attaching the SCP to the OU, all accounts within the OU will inherit the restrictions set by the SCP.

D is wrong: This option allows access only to approved AWS services by creating an SCP that allows access to only approved services and attaching it to the root OU of the organization. However, this would restrict all accounts, including those of other departments or teams within the organization. It doesn't meet the requirement of allowing each team to retain full administrative rights to its AWS account.

upvoted 2 times

🗄️ 👤 **MalonJay** 10 months, 3 weeks ago

I think Option C is wrong because the question says 'Each team also must be allowed to access only AWS services that the company approves for use'

When you deny specific services they can still access services that have not been approved.

upvoted 2 times

🗨️ 👤 **kyuhuck** 1 year ago

Selected Answer: C

Conclusion: Option C is the best solution to meet the requirements with operational efficiency and scalability. It allows teams to retain administrative rights while enforcing company-wide controls on service access through SCPs. This approach is straightforward to manage at scale, as adding or removing services from the SCP can adjust access permissions across all accounts within the OU. It directly aligns with the goal of allowing access only to approved AWS services and supports a governance model that can evolve with the organization's needs.

upvoted 3 times

🗨️ 👤 **vortegon** 1 year, 1 month ago

Selected Answer: C

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

upvoted 2 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

Selected Answer: C

C is correct: <all the accounts are in an organization in AWS Organizations> means we need scps

A and B: no mention of scps

D: SCP only denies access, not allow. Additionally, should not attach SCP to the root OU because this may inadvertently denies users' access to AWS services

upvoted 4 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

correction: D: SCP has allow statement. D perfectly fits this question

upvoted 2 times

A DevOps engineer used an AWS CloudFormation custom resource to set up AD Connector. The AWS Lambda function ran and created AD Connector, but CloudFormation is not transitioning from CREATE_IN_PROGRESS to CREATE_COMPLETE.

Which action should the engineer take to resolve this issue?

- A. Ensure the Lambda function code has exited successfully.
- B. Ensure the Lambda function code returns a response to the pre-signed URL.
- C. Ensure the Lambda function IAM role has cloudformation:UpdateStack permissions for the stack ARN.
- D. Ensure the Lambda function IAM role has ds:ConnectDirectory permissions for the AWS account.

Suggested Answer: B

Community vote distribution

B (100%)

 **haazybanj**  1 year, 2 months ago

Selected Answer: B

B. Ensure the Lambda function code returns a response to the pre-signed URL.

Explanation:

When using a custom resource in CloudFormation, the AWS Lambda function responsible for handling the resource creation should send a response to the pre-signed URL provided by CloudFormation. This response signals the completion status of the custom resource creation process to CloudFormation.

In this case, since the Lambda function successfully created the AD Connector, the engineer should ensure that the Lambda function code includes the logic to send a response to the pre-signed URL. This response should indicate the success status and any relevant data, such as the ARN or other details of the created AD Connector.


upvoted 14 times

 **dkp**  4 months, 4 weeks ago

Selected Answer: B

its B


upvoted 1 times

 **thanhnv142** 7 months, 1 week ago

B is correct: <but CloudFormation is not transitioning from CREATE_IN_PROGRESS to CREATE_COMPLETE> means ACF hasnot received a response code from Lambda

A, C and D: no mention of response code

upvoted 2 times

 **YR4591** 10 months, 2 weeks ago

Selected Answer: B

Lambda should send a cfnresponse to presign url


upvoted 1 times

 **BaburTurk** 11 months ago

B is correct

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

upvoted 3 times

 **MarDog** 1 year, 2 months ago

Selected Answer: B



It's B.

upvoted 2 times

 **OrganizedChaos25** 1 year, 3 months ago

B is correct

upvoted 2 times

  **devny** 1 year, 3 months ago

B is the right answer

upvoted 2 times

A company uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the main branch when the changes are ready for production.


The developers should not be able to push changes directly to the main branch. The company applied the `AWSCodeCommitPowerUser` managed policy to the developers' IAM role, and now these developers can push changes to the main branch directly on every repository in the AWS account.

What should the company do to restrict the developers' ability to push changes to the main branch directly?

- A. Create an additional policy to include a Deny rule for the `GitPush` and `PutFile` actions. Include a restriction for the specific repositories in the policy statement with a condition that references the main branch.
- B. Remove the IAM policy, and add an `AWSCodeCommitReadOnly` managed policy. Add an Allow rule for the `GitPush` and `PutFile` actions for the specific repositories in the policy statement with a condition that references the main branch.
- C. Modify the IAM policy. Include a Deny rule for the `GitPush` and `PutFile` actions for the specific repositories in the policy statement with a condition that references the main branch.
- D. Create an additional policy to include an Allow rule for the `GitPush` and `PutFile` actions. Include a restriction for the specific repositories in the policy statement with a condition that references the feature branches.

Suggested Answer: A

Community vote distribution



Just_Ninja Highly Voted 1 year, 8 months ago

Selected Answer: A

A is possible!

If you think C is correct, then you should know that a policy managed by AWS cannot be modified.

upvoted 15 times

jamesf Most Recent 7 months, 1 week ago

Selected Answer: A

Not C as AWS managed policy cannot be modified

upvoted 1 times

zijo 9 months ago

Selected Answer: A

AWS Managed Policies are read-only, meaning you cannot modify their contents. If you need a similar policy with slight modifications, you can copy the managed policy and create a customer-managed policy.

upvoted 2 times

dkp 11 months ago

Selected Answer: A

it s A.

upvoted 1 times

thanhnv142 1 year, 1 month ago

Selected Answer: A

A is correct: <The developers should not be able to push changes directly to the main branch> means we should deny these permissions in IAM policy. <managed polic> means we should add another policy, not modify this one.

B: <Remove the IAM policy>: this is an managed policy, cannot remove it

C: Cannot modify a managed policy. We can only create another policy

D: This option would deny committing code to every sub-branches, which is not correct

upvoted 3 times

giovanna_mag 1 year, 2 months ago

Selected Answer: A

A, AWS managed policy cannot be modified, additional policy must be attached with a DENY

upvoted 3 times

🗄️ 👤 **Blueee** 1 year, 8 months ago

Selected Answer: A

A is correct

upvoted 1 times

🗄️ 👤 **rhinozD** 1 year, 8 months ago

Selected Answer: A

AWSCodeCommitPowerUser is an AWS-managed policy.

So you need to add an additional policy to deny push to the main branch directly.

upvoted 3 times

🗄️ 👤 **Kodoma** 1 year, 9 months ago

A is correct

upvoted 1 times

🗄️ 👤 **Ryan1002** 1 year, 9 months ago

Selected Answer: A

It`s A

upvoted 2 times

🗄️ 👤 **PhuocT** 1 year, 9 months ago

Selected Answer: C

C, why we need to create an additional policy?

upvoted 2 times

🗄️ 👤 **EricZhang** 1 year, 9 months ago

You can never modify a managed policy

upvoted 5 times

🗄️ 👤 **devnv** 1 year, 10 months ago

A is correct

upvoted 3 times

A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.

A new company guideline requires a geographically isolated disaster recovery (DR) site with an RTO of 4 hours and an RPO of 15 minutes.

Which DR strategy will meet these requirements with the LEAST change to the application stack?

- A. Launch a replica environment of everything except Amazon RDS in a different Availability Zone. Create an RDS read replica in the new Availability Zone, and configure the new stack to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy.
- B. Launch a replica environment of everything except Amazon RDS in a different AWS Region. Create an RDS read replica in the new Region, and configure the new stack to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a latency routing policy.
- C. Launch a replica environment of everything except Amazon RDS in a different AWS Region. In the event of an outage, copy and restore the latest RDS snapshot from the primary Region to the DR Region. Adjust the Route 53 record set to point to the ALB in the DR Region.
- D. Launch a replica environment of everything except Amazon RDS in a different AWS Region. Create an RDS read replica in the new Region, and configure the new environment to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy. In the event of an outage, promote the read replica to primary.

Suggested Answer: D

Community vote distribution

D (89%)

11%

 **YR4591** Highly Voted 10 months, 2 weeks ago

Selected Answer: D

D is correct. Failover policy will route traffic to the ALB in the backup region.
upvoted 5 times

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

Selected Answer: D

D is correct. My answer was C, but GPT told me this:
This process does not guarantee a quick recovery within the 4-hour RTO. Restoring from a snapshot would take significant time, and it doesn't provide continuous replication of data to minimize RPO. This option does not meet the RTO requirement because the recovery process (snapshot restore) will take too long. It also doesn't address continuous data replication, leading to a higher potential RPO.
upvoted 1 times

 **dkp** 4 months, 4 weeks ago

Selected Answer: D

geographically isolated location applies to option D
upvoted 2 times

 **WhyIronMan** 5 months, 1 week ago

Selected Answer: D

Answer is D.

A. It did not cover the whole scenario. there is a need to promote the read replica otherwise the application in the region will not be able to write in that rds. Also, another Region means geographically isolated, while other Az will not solve the problem.

Details are everything during an investigation...
upvoted 3 times

 **WhyIronMan** 5 months, 1 week ago

Answer is D.

A. It did not cover the whole scenario. there is a need to promote the read replica otherwise the application in the region will not be able to write in that rds. Also, another Region means geographically isolated, while other Az will not solve the problem.

Details are everything during an investigation...

upvoted 1 times

🗳️ 👤 **jojom19980** 6 months, 3 weeks ago

Selected Answer: A

they mentioned geographically not regional , the cost will be more if we make it regional so we can go with the AZ DR

upvoted 2 times

🗳️ 👤 **WhyIronMan** 5 months, 1 week ago

A did not cover the whole scenario. there is a need to promote the read replica otherwise the application in the region will not be able to write in that rds. Also, another Region means geographically isolated.

Details are everything during an investigation...

upvoted 1 times

🗳️ 👤 **thanhnv142** 7 months, 1 week ago

Selected Answer: D

D is correct: < configured Amazon Route 53> and <requires a geographically isolated disaster recovery (DR) site> means fail-over routing and the DR site should be in another region

A: <Amazon RDS in a different Availability Zone>: We need to setup the DB in a different region, not in a different AZ, which is still in the same region

B and C: no mention of fail-over

upvoted 3 times

🗳️ 👤 **sarlos** 8 months, 1 week ago

D is the answer

upvoted 2 times

🗳️ 👤 **HugoFM** 9 months, 2 weeks ago

What about C? I have an RTO of 4 hours I mean we got time we dont need a read replica we could take time to restore from a snapshot

upvoted 3 times

🗳️ 👤 **davdan99** 8 months ago

Restore Time vs. RTO: While 4 hours might seem like a sufficient window for restoring a snapshot, the actual restore time can vary depending on several factors:

Snapshot size: Larger snapshots take longer to restore.

RDS instance type: High-performance instance types can handle restorations faster.

Network bandwidth: Sufficient bandwidth is crucial for speedy data transfer during restoration.

RDS engine version: Newer versions might have optimized restore processes.

In practice, restoring a large RDS snapshot, especially across Regions, could easily take more than 15 minutes, potentially exceeding the RPO and resulting in data loss.

upvoted 2 times

🗳️ 👤 **haazybanj** 1 year, 2 months ago

Selected Answer: D

D is right

upvoted 4 times

🗳️ 👤 **Kodoma** 1 year, 3 months ago

D is true

upvoted 2 times

🗳️ 👤 **devnv** 1 year, 3 months ago

D is correct

upvoted 1 times

A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS for Oracle DB instance and Amazon DynamoDB. There are separate environments for development, testing, and production.

What is the MOST secure and flexible way to obtain password credentials during deployment?

- A. Retrieve an access key from an AWS Systems Manager SecureString parameter to access AWS services. Retrieve the database credentials from a Systems Manager SecureString parameter.
- B. Launch the EC2 instances with an EC2 IAM role to access AWS services. Retrieve the database credentials from AWS Secrets Manager.
- C. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS services. Retrieve the database credentials from a Systems Manager SecureString parameter.
- D. Launch the EC2 instances with an EC2 IAM role to access AWS services. Store the database passwords in an encrypted config file with the application artifacts.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ **jamesf** 7 months, 2 weeks ago

Selected Answer: B

Keywords: MOST secure
upvoted 2 times

🗳️ **zijo** 9 months ago

Selected Answer: B

This step is important for applications running on EC2 instances to retrieve passwords from AWS Secrets Manager. Create an IAM role with the necessary permissions to access AWS Secrets Manager. Attach this IAM role to your EC2 instance.
upvoted 2 times

🗳️ **c3518fc** 10 months, 3 weeks ago

Selected Answer: B

The most secure and flexible way to obtain password credentials during deployment in the given scenario is to use AWS Secrets Manager. AWS Secrets Manager is a service that allows you to securely store, retrieve, and rotate credentials, such as passwords, API keys, and other sensitive data.
upvoted 3 times

🗳️ **dkp** 11 months ago

Selected Answer: B

B seems more relevant
upvoted 2 times

🗳️ **WhyIronMan** 11 months, 2 weeks ago

Selected Answer: B

B. EC2 Role + Secrets Manager
upvoted 2 times

🗳️ **thanhnv142** 1 year, 1 month ago

Selected Answer: B

B is correct: <obtain password credentials> means we should consider AWS SSM and secret manager. However, <the MOST secure > means we should opt for secret manager, which is more costly but more secure
A, C and D: no mention of secret manager
upvoted 4 times

🗳️ **sarlos** 1 year, 2 months ago

why not A?

upvoted 1 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

<obtain password credentials> means we should consider AWS SSM and secret manager. However, <the MOST secure > means we should opt for secret manager, which is more costly but more secure

upvoted 3 times

🗨️ 👤 **davdan99** 1 year, 2 months ago

We are not storing access keys for EC2 instances, instead we are using instance profile for that it is the best practice, and for database credentials it is correct to use Secret manager, it is more integrated with RDS, and other database services within AWS.

upvoted 1 times

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

Selected Answer: B

I vote B

upvoted 2 times

🗨️ 👤 **Snape** 1 year, 7 months ago

Selected Answer: B

No Brainer

upvoted 3 times

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

Selected Answer: B

Most secure is B

upvoted 4 times

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

Selected Answer: B

Option B is the right answer.

upvoted 2 times

🗨️ 👤 **devnv** 1 year, 10 months ago

B sounds the right answer

upvoted 1 times


The security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.

What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

- A. Create an Amazon EventBridge rule for the CloudTrail StopLogging event. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called. Add the Lambda function ARN as a target to the EventBridge rule.
- B. Deploy the AWS-managed CloudTrail-enabled AWS Config rule, set with a periodic interval of 1 hour. Create an Amazon EventBridge rule for AWS Config rules compliance change. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called. Add the Lambda function ARN as a target to the EventBridge rule.
- C. Create an Amazon EventBridge rule for a scheduled event every 5 minutes. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS account. Add the Lambda function ARN as a target to the EventBridge rule.
- D. Launch a t2.nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current account. If the CloudTrail trail is disabled, have the script re-enable the trail.

Suggested Answer: A

Community vote distribution



rhinozD Highly Voted 1 year, 2 months ago

Selected Answer: A

A.

old but gold link:

<https://aws.amazon.com/blogs/mt/monitor-changes-and-auto-enable-logging-in-aws-cloudtrail/>

upvoted 15 times

daburahjail 11 months, 3 weeks ago

Good job, buddy

upvoted 2 times

c3518fc Most Recent 4 months, 3 weeks ago

Selected Answer: A

This solution ensures the least amount of downtime for CloudTrail log deliveries when auto-remediating CloudTrail being turned off. Here's why:

Event-Driven Automation: By creating an Amazon EventBridge rule for the CloudTrail StopLogging event, the remediation process is triggered immediately when CloudTrail logging is stopped, minimizing the downtime.

Targeted Remediation: The Lambda function uses the AWS SDK to call StartLogging on the specific CloudTrail trail ARN where the StopLogging event occurred. This targeted approach ensures that logging is re-enabled for the affected trail without impacting other trails or introducing unnecessary overhead.

Low Latency: EventBridge rules and Lambda functions are designed to be highly responsive, ensuring that the remediation action is initiated with minimal delay after the StopLogging event occurs.

upvoted 2 times

WhyIronMan 5 months, 1 week ago

Selected Answer: A

A. is correct

Details are everything during an investigation

upvoted 1 times

thanhnv142 7 months, 1 week ago

A is correct: <The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off> means we should turn on it again if we detect that it is turn-off. AWS config rule or Eventbridge would be considered. < the LEAST amount of downtime> means we should choose A because this minimizes downtime

B: this option utilizes an AWS config rule, which is good. But it sets the rule with a periodic interval of 1 hours, which would introduce a lot of

downtime

C: this option utilizes evenbridge, but the event to trigger eventbridge is undetermined

D: Should not use a custom script to do the task

upvoted 2 times

🗳️ 👤 **HugoFM** 9 months, 2 weeks ago

Selected Answer: A

A its quicker and the solution is asking for the leeast amount of downtime

upvoted 2 times

🗳️ 👤 **YR4591** 10 months, 2 weeks ago

Selected Answer: A

"LEAST amount of downtime" = A

cloudwatch event is near real time. Al the other options are not.

upvoted 1 times

🗳️ 👤 **RVivek** 1 year ago

Selected Answer: A

Both A and B will work. However the question mentions leatst Cloutrial down time. Option A is correct beacuse the remiation is triggred immeiately .

Option B can be delayed a it uns once in ever hour

upvoted 1 times

🗳️ 👤 **Seoyong** 1 year ago

I don't think stoplogging is CloudTrail being turned off.

You can stop logging anytime - <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-turning-off-logging.html>

But it doesn't means CloudTrail being turned off

upvoted 1 times

🗳️ 👤 **Seoyong** 1 year ago

Selected Answer: B

cloudtrail-enabled rule will check CloudTrail being turned off.

upvoted 1 times

🗳️ 👤 **WhyIronMan** 5 months, 1 week ago

Please notice that the question says "LEAST amount of downtime" while B is possible it says "set with a periodic interval of 1 hour." which can basically take 1h to enable CloudTrail again with causes a lot of downtime

upvoted 2 times

🗳️ 👤 **OrganizedChaos25** 1 year, 3 months ago

I got A as my answer

upvoted 2 times

🗳️ 👤 **Mail1964** 1 year, 3 months ago

The requirement is - What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries? For me that means reacting to AWS events. AWS config rule with 1 hr schedule does not meet the criteria in my opinion.

upvoted 2 times

🗳️ 👤 **2pk** 1 year, 3 months ago

Selected Answer: A

Answer A,

This solution is the most appropriate as it listens to the StopLogging event and automatically starts logging immediately. This approach eliminates the need to wait for a scheduled interval, thereby reducing the amount of downtime and ensuring the security team can detect security issues in real-time.

Option B is incorrect as it uses AWS Config rules to detect CloudTrail stoppage, which might not be an immediate solution to this issue

upvoted 3 times

🗳️ 👤 **devnv** 1 year, 3 months ago

A is the right answer

upvoted 2 times

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

Selected Answer: A

A is correct.

upvoted 2 times

  **Jeanphi72** 1 year, 4 months ago

Selected Answer: B

B for me

upvoted 1 times

A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "codeartifact:DescribePackageVersion",
        "codeartifact:DescribeRepository",
        "codeartifact:GetPackageVersionReadme",
        "codeartifact:GetRepositoryEndpoint",
        "codeartifact:ListPackageVersionAssets",
        "codeartifact:ListPackageVersionDependencies",
        "codeartifact:ListPackageVersions",
        "codeartifact:Listpackages",
        "codeartifact:ReadFromRepository"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Principal": "*"
    },
    {
      "Condition": {
        "StringEquals": {
          "aws:PrincipalOrgID": [
            "o-xxxxxxxxxxxx"
          ]
        }
      }
    }
  ]
}
```

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC. Because of compliance requirements, the VPC has no internet connectivity.


The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec.yaml file. However, the development team cannot download the Python library from the repository.

Which combination of steps should a DevOps engineer take so that the development team can use CodeArtifact? (Choose two.)

- A. Create an Amazon S3 gateway endpoint. Update the route tables for the subnets that are running the CodeBuild job.
- B. Update the repository policy's Principal statement to include the ARN of the role that the CodeBuild project uses.
- C. Share the CodeArtifact repository with the organization by using AWS Resource Access Manager (AWS RAM).
- D. Update the role that the CodeBuild project uses so that the role has sufficient permissions to use the CodeArtifact repository.
- E. Specify the account that hosts the repository as the delegated administrator for CodeArtifact in the organization.

Suggested Answer: BD

Community vote distribution



TroyMcLure Highly Voted 1 year, 9 months ago

Selected Answer: AD

I guess the answer is AD because of this:

"AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."

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

With no internet connectivity, a gateway endpoint becomes necessary to access S3.

upvoted 13 times

  **Arnaud92** 1 year, 9 months ago

<https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>

It clearly state that you need to create a S3 endpoint to use codeartifact in a private network.

upvoted 8 times

  **vortegon** 1 year, 1 month ago

An Amazon S3 endpoint is not needed when using Python or Swift package formats.

upvoted 5 times


  **syh_rapha** 8 months ago

When this question was created, there was no exception for Python and Swift packages. You can check this using the Wayback machine:

<https://web.archive.org/web/20230521063821/https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>



Considering that it's very common to have outdated questions in the exam, I'd say this is one those cases. So yeah, I'll also go with AD (also because B is not needed since the repository policy is already allowing the entire org).

upvoted 2 times

  **RVivek** 1 year, 5 months ago

A- incorrect because the question says Devops engineers created VPC endpoints for CodeArtifact

upvoted 2 times

  **RVivek** 1 year, 5 months ago

AD even though Devops engineer created a CodeArtifact still a S3 end point is required

upvoted 2 times

  **Venki_dev** 8 months, 3 weeks ago

note here says "An Amazon S3 endpoint is not needed when using Python or Swift package formats."

<https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>

upvoted 1 times

  **Jowblow**  1 year, 10 months ago

Selected Answer: AD

Codeartifact uses s3 gateway endpoints to store packages. The key word here are no internet access.

upvoted 6 times

  **79f3aa3**  2 weeks ago



Selected Answer: BD

Two things that are clearly mentioned in the question:

1. Python library has already been stored in the CodeArtifact repository
2. VPC endpoint for CodeArtifact also have been configured

The only thing left are the permissions to access CodeArtifact, which mean B and D.

upvoted 1 times

  **spring21** 3 months, 1 week ago

Selected Answer: AD

AWS CodeArtifact stores artifact data and metadata in Amazon S3. To pull packages from CodeArtifact, you need to create an Amazon S3 gateway endpoint. You can use the aws ec2 create-vpc-endpoint AWS CLI command to create the endpoint.

upvoted 1 times

  **steli0** 3 months, 2 weeks ago

Selected Answer: BD

I vote for BD even though it's not so clear if both repo and CodeBuild are in the same org. Moreover, S3 GW endpoint auto-creates the routes with prefix-lists at your table. <https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html#create-gateway-endpoint-s3>

Nevertheless as mentioned here in AWS documentation that other users posted S3 .

I hope this question is removed from the exam.

upvoted 1 times

  **[Removed]** 6 months, 3 weeks ago

Selected Answer: AD

AD for me

upvoted 1 times

  **iulian0585** 7 months ago

Selected Answer: AD



<https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>

upvoted 2 times

  **auxwww** 7 months, 2 weeks ago

B - doesn't make any sense because aws:PrincipalOrgID condition key in repo policy already allows any principal within the org to access the repo

upvoted 2 times

  **Venki_dev** 8 months, 3 weeks ago



Selected Answer: BD

BD

note here says "An Amazon S3 endpoint is not needed when using Python or Swift package formats."

<https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>

upvoted 1 times

  **zijo** 9 months ago

Selected Answer: CD

C is needed

if the codeartifact and codebuild are in different organization accounts, AWS RAM is a service that allows you to share AWS resources with other AWS accounts within your organization. AWS RAM can be used to share CodeArtifact resources across different accounts.

A is not needed

you do not need an S3 gateway as a VPC endpoint specifically for using AWS CodeArtifact with Python packages. AWS CodeArtifact itself manages the storage and retrieval of packages, and it uses its own service endpoints for these operations.

D is needed for

Ensure the IAM role used by CodeBuild has permissions to access CodeArtifact

B is not needed

Here it is not required because the CodeArtifact policy has Principal as *

upvoted 3 times

  **that1guy** 10 months ago

Selected Answer: CD

C and D

A - S3 gateway endpoint is not required for Python: <https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>

B - Principal is already "*".

upvoted 1 times

  **vn_thanhtung** 9 months, 2 weeks ago

Pls Read link <https://docs.aws.amazon.com/ram/latest/userguide/shareable.html>

upvoted 1 times

  **seetpt** 10 months, 1 week ago

Selected Answer: AD

AD because Principal is already "*".

upvoted 2 times

  **xdkonorek2** 10 months, 2 weeks ago

Selected Answer: BD

as for A: "To pull packages from CodeArtifact, you must create a gateway endpoint for Amazon S3." but... "Note - An Amazon S3 endpoint is not needed when using Python or Swift package formats."

<https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>

upvoted 2 times

  **c3518fc** 10 months, 3 weeks ago

Selected Answer: BD



The issue here is policy update as the developers have already enabled VPC endpoint (CodeArtifact uses Amazon Simple Storage Service (Amazon S3) to store package assets. To pull packages from CodeArtifact, you must create a gateway endpoint for Amazon S3. When your build or deployment process downloads packages from CodeArtifact, it must access CodeArtifact to get package metadata and Amazon S3 to download package assets (for example, Maven .jar files).

Note

An Amazon S3 endpoint is not needed when using Python or Swift package formats.

To create the Amazon S3 gateway endpoint for CodeArtifact, use the Amazon EC2 create-vpc-endpoint AWS CLI command. When you create the endpoint, you must select the route tables for your VPC. For more information, see Gateway VPC Endpoints in the Amazon Virtual Private Cloud User Guide.)

upvoted 4 times

  **c3518fc** 10 months, 3 weeks ago

<https://docs.aws.amazon.com/codeartifact/latest/ug/create-s3-gateway-endpoint.html>

upvoted 1 times

  **dkp** 11 months ago

Selected Answer: BD

ANS B&D

CodeArtifact uses Amazon Simple Storage Service (Amazon S3) to store package assets. To pull packages from CodeArtifact, you must create a gateway endpoint for Amazon S3. When your build or deployment process downloads packages from CodeArtifact, it must access CodeArtifact to get package metadata and Amazon S3 to download package assets (for example, Maven .jar files).

Note

An Amazon S3 endpoint is not needed when using Python or Swift package formats.

upvoted 4 times

  **WhyIronMan** 11 months, 2 weeks ago

Selected Answer: AD

A,D are correct

upvoted 3 times

  **kyuhuck** 1 year ago

Selected Answer: AD

'ad' correct = 'AWS codeartiface' operates in multiple availability zones and stores artiface data and metadata in amazon s3 and amazon dynamoDB your encrypted data is redundantly stored across myltiple facilities and multiple devices in each facility, marking it highly available and highly durable...

upvoted 3 times

A company uses a series of individual Amazon CloudFormation templates to deploy its multi-Region applications. These templates must be deployed in a specific order. The company is making more changes to the templates than previously expected and wants to deploy new templates more efficiently. Additionally, the data engineering team must be notified of all changes to the templates.

What should the company do to accomplish these goals?



- A. Create an AWS Lambda function to deploy the CloudFormation templates in the required order. Use stack policies to alert the data engineering team.
- B. Host the CloudFormation templates in Amazon S3. Use Amazon S3 events to directly trigger CloudFormation updates and Amazon SNS notifications.
- C. Implement CloudFormation StackSets and use drift detection to trigger update alerts to the data engineering team.
- D. Leverage CloudFormation nested stacks and stack sets for deployments. Use Amazon SNS to notify the data engineering team.

Suggested Answer: B

Community vote distribution

D (93%)

7%

  **emupsx1** Highly Voted 1 year, 1 month ago

The answer is D because:

A few hours ago, I just finished the DOP-C02 exam.

My score is 1000 points.

This question has come up, I choose D.

upvoted 13 times

  **BaburTurk** 1 year ago

pics or it did not happen, troll bot account

upvoted 9 times

  **c3518fc** Highly Voted 4 months, 3 weeks ago

Selected Answer: D

Here's why this solution is the best approach:

Nested Stacks: CloudFormation nested stacks allow you to break down complex templates into smaller, more manageable templates. You can create a root stack that references and manages multiple nested stacks. This approach simplifies the management and deployment of multiple interdependent templates in the correct order.

StackSets: CloudFormation StackSets allow you to create, update, or delete stacks across multiple AWS accounts and regions with a single operation. This addresses the requirement of deploying applications across multiple regions efficiently.

Amazon SNS: Amazon Simple Notification Service (SNS) can be used to send notifications to the data engineering team whenever changes are made to the CloudFormation templates or stacks.

upvoted 5 times

  **dkp** Most Recent 4 months, 4 weeks ago

Selected Answer: D

should be D



upvoted 1 times

  **WhyIronMan** 5 months, 1 week ago

Selected Answer: D

Answer is D.

upvoted 1 times

  **jojom19980** 6 months, 3 weeks ago

Selected Answer: D

the nested for order :

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-nested-stacks.html>

upvoted 2 times

🗨️ 👤 **thanhnv142** 7 months, 1 week ago

D is correct: <uses a series of individual Amazon CloudFormation templates to deploy its multi-Region applications> and <wants to deploy new templates more efficiently> mean stacksets, which is template for multiple regions. <the data engineering team must be notified of all changes> means SNS

A and B: no mention of stacksets

C: no mention of SNS

upvoted 3 times

🗨️ 👤 **sarlos** 8 months, 1 week ago

D is the answer

upvoted 1 times

🗨️ 👤 **YR4591** 10 months, 2 weeks ago

Selected Answer: D

It's D.

C is not correct since according to this link:

<https://aws.amazon.com/blogs/mt/implementing-an-alarm-to-automatically-detect-drift-in-aws-cloudformation-stacks/>

We need AWS config rule to detect drifts and to send event. There is no build in solution to notify drift detection like mentioned in C,

upvoted 1 times

🗨️ 👤 **daburahjail** 11 months, 3 weeks ago

Selected Answer: D

C is for notifying changes on what has been deployed by Cloud Formation

D is for notifying changes made on the Cloud Formation template (the recipe) itself

upvoted 1 times

🗨️ 👤 **beanxyz** 1 year ago

Selected Answer: D

C didn't mention how to deploy in a specific order.

D mentioned nested stack, which you can configure the dependency ordering

upvoted 1 times

🗨️ 👤 **OrganizedChaos25** 1 year, 3 months ago

I got D as my answer

upvoted 1 times

🗨️ 👤 **devnv** 1 year, 3 months ago

D is correct

upvoted 1 times

🗨️ 👤 **PhuocT** 1 year, 3 months ago

Selected Answer: D

it' s D, I think.

upvoted 1 times

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

Selected Answer: D

D is correct.

upvoted 1 times

🗨️ 👤 **Sazeka** 1 year, 4 months ago

Selected Answer: D

The correct solution is D:

The solution works as follows:

AWS Config triggers the evaluation when any resource that matches the rule's scope (currently set to "AWS::CloudFormation::Stack") changes in configuration and at the frequency ("MaximumExecutionFrequency" parameter) that you specify at the time of this solution deployment.



EventBridge receives the events from AWS Config, applies the EventBridge rule to match the compliance change event, and transforms the input (customize the text) as defined in the "InputTransformer" template.

The chosen customer-managed KMS Key is then accessed to encrypt the notification.

The encrypted notification is published to the target SNS topic.

The endpoints subscribed to this topic start receiving the published messages.

upvoted 1 times

  **Jeanphi72** 1 year, 4 months ago

Selected Answer: C

I think C: <https://aws.amazon.com/blogs/mt/implementing-an-alarm-to-automatically-detect-drift-in-aws-cloudformation-stacks/>

upvoted 1 times

A DevOps engineer has implemented a CI/CD pipeline to deploy an AWS CloudFormation template that provisions a web application. The web application consists of an Application Load Balancer (ALB), a target group, a launch template that uses an Amazon Linux 2 AMI, an Auto Scaling group of Amazon EC2 instances, a security group, and an Amazon RDS for MySQL database. The launch template includes user data that specifies a script to install and start the application.

The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The CI/CD pipeline has deployed a new version of the template. However, the health checks on the ALB are now failing. The health checks have marked all targets as unhealthy.

During investigation, the DevOps engineer notices that the CloudFormation stack has a status of UPDATE_COMPLETE. However, when the DevOps engineer connects to one of the EC2 instances and checks /var/log/messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error.

How can the DevOps engineer ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running?

- A. Use the cfn-signal helper script to signal success or failure to CloudFormation. Use the WaitOnResourceSignals update policy within the CloudFormation template. Set an appropriate timeout for the update policy.
- B. Create an Amazon CloudWatch alarm for the UnhealthyHostCount metric. Include an appropriate alarm threshold for the target group. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation.
- C. Create a lifecycle hook on the Auto Scaling group by using the AWS::AutoScaling::LifecycleHook resource. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation. Set an appropriate timeout on the lifecycle hook.
- D. Use the Amazon CloudWatch agent to stream the cloud-init logs. Create a subscription filter that includes an AWS Lambda function with an appropriate invocation timeout. Configure the Lambda function to use the SignalResource API operation to signal success or failure to CloudFormation.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ **Certified101** Highly Voted 1 year, 7 months ago

Selected Answer: A

A is correct

upvoted 6 times

🗳️ **zijo** Most Recent 9 months ago

Selected Answer: A

To ensure that the CloudFormation deployment fails if the user data script does not successfully finish, you can use a combination of AWS CloudFormation's CreationPolicy, cfn-signal, and wait condition resources. These mechanisms can signal CloudFormation about the success or failure of the instance creation process, including the execution of user data scripts.

upvoted 3 times

🗳️ **dkp** 11 months ago

Selected Answer: A

A is correct

upvoted 2 times

🗳️ **thanhnv142** 1 year, 1 month ago

A is correct: <ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running> means we need cfn-signal
B, C and D: no mention of cfn-signal

upvoted 3 times

🗳️ **sarlos** 1 year, 2 months ago

yes A.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-signal.html>

upvoted 3 times

🗨️ 👤 **svjl** 1 year, 3 months ago

The instance is running, and the logs are available. The configuration happens inside the instance by the userdata. How A is correct if the issue is beyond CF?

upvoted 1 times

🗨️ 👤 **svjl** 1 year, 3 months ago

Ok now make sense: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-signal.html>

upvoted 2 times

🗨️ 👤 **OrganizedChaos25** 1 year, 9 months ago

A is correct

upvoted 1 times

🗨️ 👤 **devnv** 1 year, 10 months ago

A is the right answer

upvoted 1 times

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

Selected Answer: A

Agree with A

upvoted 2 times

A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently, the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive. Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically for the application.

To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed.

Which solution will meet these requirements?

- A. Create an Amazon EventBridge rule to send notifications to the security team whenever a user logs in to an EC2 instance. Use EC2 Instance Connect to log in to the instances. Deploy Auto Scaling groups by using AWS CloudFormation. Use the cfn-init helper script to deploy appropriate VPC routes for external access. Rebuild the custom AMI so that the custom AMI includes AWS Systems Manager Agent.
- B. Deploy a NAT gateway and a bastion host that has internet access. Create a security group that allows incoming traffic on all the EC2 instances from the bastion host. Install AWS Systems Manager Agent on all the EC2 instances. Use Auto Scaling group lifecycle hooks for monitoring and auditing access. Use Systems Manager Session Manager to log in to the instances. Send logs to a log group in Amazon CloudWatch Logs. Export data to Amazon S3 for auditing. Send notifications to the security team by using S3 event notifications.
- C. Use EC2 Image Builder to rebuild the custom AMI. Include the most recent version of AWS Systems Manager Agent in the image. Configure the Auto Scaling group to attach the AmazonSSMManagedInstanceCore role to all the EC2 instances. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.
- D. Use AWS Systems Manager Automation to build Systems Manager Agent into the custom AMI. Configure AWS Config to attach an SCP to the root organization account to allow the EC2 instances to connect to Systems Manager. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.

Suggested Answer: C

Community vote distribution

C (100%)

 **Blueee**  1 year, 8 months ago

Selected Answer: C

C and D are left over choice due to no internet access for EC2

C is correct

By using EC2 Image Builder to rebuild the custom AMI and including the most recent version of AWS Systems Manager Agent in the image, you can configure the Auto Scaling group to attach the AmazonSSMManagedInstanceCore role to all the EC2 instances. This allows you to use Systems Manager Session Manager to log in to the instances. You can enable logging of session details to Amazon S3 and create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic2 upvoted 7 times

 **thanhnv142**  1 year, 1 month ago

C is correct: <The company needs to monitor the application and consolidate access to the application> means using SSM. We should install SSM agent on all EC2 instances. <The EC2 instances run a custom AMI that is built specifically for the application> means we should rebuild the image and integrate agent into the AMI. To rebuild, the best option is EC2 image builder. <The company's security team must receive a notification whenever the instances are accessed.> means SNS

A: <Rebuild the custom AMI so that the custom AMI includes AWS Systems Manager Agent.>: no mention of using EC2 image builder and SNS

B: no mention of integrating SSM agents into the AMI and we cannot just send S3 noti to users <Send notifications to the security team by using S3 event notifications.>

D: no me ntion of using EC2 image builder to rebuild the AMI.

upvoted 6 times

🗨️ 👤 **Saudis** Most Recent 3 months, 3 weeks ago

Ans is C because The keyword is access must be automated and controlled centrally
upvoted 1 times

🗨️ 👤 **jamesf** 7 months, 1 week ago

Selected Answer: C

- AWS Systems Manager Agent
- Systems Manager Session Manager for login the instances
- enable logging of session details to s3
- s3 event notification to SNS.

upvoted 1 times

🗨️ 👤 **dkp** 11 months ago

Selected Answer: C

C is correct
upvoted 1 times

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

Selected Answer: C

C

Option C offers a well-architected approach to addressing the requirements, providing both centralized access and logging, and automated login to EC2 instances for system administrators. Additionally, it ensures that the security team receives notifications for auditing and monitoring purposes.

upvoted 2 times

🗨️ 👤 **PhuocT** 1 year, 9 months ago

D is not a good option for the following reasons:

1. AWS Systems Manager Automation is not the ideal choice for building a custom AMI. Instead, EC2 Image Builder, as stated in option C, is an AWS service designed for building, testing, and maintaining Golden Amazon Machine Images (AMIs), making it a suitable choice for both building and managing custom AMIs.

2. The option D suggests attaching an SCP (Service Control Policy) to the root organization to allow EC2 instances to connect to Systems Manager. This approach is incorrect because SCPs are used to define permissions on an organizational level, rather than allowing specific access between resources like EC2 instances and Systems Manager. Attaching the AmazonSSMManagedInstanceCore role to EC2 instances as mentioned in option C is the correct method, which allows instances to communicate with Systems Manager.

upvoted 3 times

🗨️ 👤 **2pk** 1 year, 10 months ago

if someone know why D is not correct , pls post
upvoted 1 times

🗨️ 👤 **MarDog** 1 year, 8 months ago

Because I don't think AWS Config can be used to attach an SCP.
upvoted 2 times

A company uses Amazon S3 to store proprietary information. The development team creates buckets for new projects on a daily basis. The security team wants to ensure that all existing and future buckets have encryption, logging, and versioning enabled. Additionally, no buckets should ever be publicly read or write accessible.

What should a DevOps engineer do to meet these requirements?

- A. Enable AWS CloudTrail and configure automatic remediation using AWS Lambda.
- B. Enable AWS Config rules and configure automatic remediation using AWS Systems Manager documents.
- C. Enable AWS Trusted Advisor and configure automatic remediation using Amazon EventBridge.
- D. Enable AWS Systems Manager and configure automatic remediation using Systems Manager documents.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **dkp** 4 months, 4 weeks ago

Selected Answer: B

AWS config to remediate non-compliance
upvoted 2 times

🗳️ 👤 **thanhnv142** 7 months, 1 week ago

Selected Answer: B

B is correct: <wants to ensure that all existing and future buckets have encryption, logging, and versioning enabled> means we need aws config.
A, C and D: no mention of AWS config
upvoted 4 times

🗳️ 👤 **yuliaqwerty** 8 months ago

Answer B

upvoted 1 times

🗳️ 👤 **sarlos** 8 months, 1 week ago

yes B is correct

upvoted 1 times

🗳️ 👤 **Arnaud92** 1 year, 3 months ago

Selected Answer: B

AWS Config allows you to remediate noncompliant resources that are evaluated by AWS Config Rules. AWS Config applies remediation using AWS Systems Manager Automation documents.

see <https://docs.aws.amazon.com/config/latest/developerguide/remediation.html>

upvoted 4 times

🗳️ 👤 **OrganizedChaos25** 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 3 times

A DevOps engineer is researching the least expensive way to implement an image batch processing cluster on AWS. The application cannot run in Docker containers and must run on Amazon EC2. The batch job stores checkpoint data on an NFS volume and can tolerate interruptions. Configuring the cluster software from a generic EC2 Linux image takes 30 minutes.

What is the MOST cost-effective solution?

- A. Use Amazon EFS for checkpoint data. To complete the job, use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.
- B. Use GlusterFS on EC2 instances for checkpoint data. To run the batch job, configure EC2 instances manually. When the job completes, shut down the instances manually.
- C. Use Amazon EFS for checkpoint data. Use EC2 Fleet to launch EC2 Spot Instances, and utilize user data to configure the EC2 Linux instance on startup.
- D. Use Amazon EFS for checkpoint data. Use EC2 Fleet to launch EC2 Spot Instances. Create a custom AMI for the cluster and use the latest AMI when creating instances.

Suggested Answer: C

Community vote distribution

D (88%)

8%

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

Selected Answer: D

D is more suitable, as it says to avoid 30min launch time.

upvoted 7 times

🗳️ 👤 **Mail1964** 1 year, 9 months ago

I assume you are saying D over C as D will make the EC2 instances operational quicker, while C would require 30 minutes to install the software before it can start to be used. resulting in it being more cost effective.

upvoted 3 times

🗳️ 👤 **jamesf** Most Recent 7 months, 2 weeks ago

Selected Answer: D

keywords: MOST cost-effective, a generic EC2 Linux image takes 30 minutes.

Mean take 30mins or longer time for EC2 booting and will cost more.

upvoted 2 times

🗳️ 👤 **WhyIronMan** 11 months, 2 weeks ago

Selected Answer: D

D is the correct answer.

Make the calculations 30 min of bootstrapping when you have multiple scale actions is a lot of time idle when you have many instances, so the money spent during a lot of spot request that waste time bootstrapping is larger than keeping a single AMI.

Details are everything during an investigation...

upvoted 3 times

🗳️ 👤 **Diego1414** 1 year ago

Selected Answer: C

Answer is C.

C is cheaper than D

upvoted 2 times

🗳️ 👤 **HayLLiHuK** 11 months, 1 week ago

"utilize user data to configure the EC2 Linux instance on startup" - it takes an addition time to configure an instance.

it's better to use a custom AMI and have everything preinstalled

upvoted 3 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

C is correct: <can tolerate interruptions> means EC2 spot instances.

A and B: no mention of spot instances

D: Create a custom AMI for the cluster and use the latest AMI when creating instances: this incurs more cost than option C, which incurs no cost for the configuration step

upvoted 3 times

  **vmahilevskyi** 11 months, 4 weeks ago

As for me, extra 30 minutes for each EC2 launch seems like an extra cost comparing to one-time built AMI.

So D looks cheaper than C

upvoted 3 times

  **zain1258** 1 year, 3 months ago

Selected Answer: D

It's D

upvoted 2 times

  **Cloud_noob** 1 year, 4 months ago

Selected Answer: D

D most suitable. why would you be willing to wait for 30 minutes for software installation?

upvoted 2 times

  **FEEREWMWKA** 1 year, 6 months ago

D - Spot Instances are cheaper than Ec2 and the workload can tolerate interruptions. Also Custom AMI removes the need for 30 mins configuration which takes a resource that will need to be paid

upvoted 4 times

  **thanhnv142** 1 year, 1 month ago

I think you may be wrong. D - creating custom AMI costs you a lot. You would need to buy services from EC2 image builder to rebuild and pay for the one who rebuild the custom AMI. Meanwhile, configuring manually (option C) costs you nothing but time. Of course, in corporate environment, time is money, but you would cost way less than option D - which suggests cost from creating and maintaining the AMI image (latest AMI image)

upvoted 2 times

  **ogwu2000** 1 year, 7 months ago

C is the answer. Because it specifically mentioned EC2 Linux instance as requested in the question. D only mentioned creating an AMI but is it Linux AMI?

upvoted 3 times

  **2pk** 1 year, 4 months ago

Read the question properly. :D It says Linux EC2 instances..



upvoted 2 times

  **OrganizedChaos25** 1 year, 9 months ago

Selected Answer: D



D is correct

upvoted 4 times

  **devnv** 1 year, 9 months ago

D is the right answer



upvoted 3 times

  **meisme** 1 year, 10 months ago

Selected Answer: D

D is correct

upvoted 2 times

  **Jeanphi72** 1 year, 10 months ago

Selected Answer: B

B because: " Configuring the cluster software from a generic EC2 Linux image takes 30 minutes.

"

upvoted 1 times


A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer, which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.

Which solution will meet these requirements with MINIMAL changes to the application?

- A. Introduce changes as a separate environment parallel to the existing one. Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.
- B. Introduce changes as a separate environment parallel to the existing one. Update the application's DNS alias records to point to the new environment.
- C. Introduce changes as a separate target group behind the existing Application Load Balancer. Configure API Gateway to route user traffic to the new target group in steps.
- D. Introduce changes as a separate target group behind the existing Application Load Balancer. Configure API Gateway to route all traffic to the Application Load Balancer, which then sends the traffic to the new target group.

Suggested Answer: C

Community vote distribution



haazybanj Highly Voted 1 year, 7 months ago

Selected Answer: A

Option A is also a valid approach that can meet the requirements with MINIMAL changes to the application.

In Option A, the changes are introduced as a separate environment parallel to the existing one. This new environment can be used to deploy the new version of the application. By configuring API Gateway to use a canary release deployment, a small subset of user traffic is directed to the new environment, while the majority of traffic continues to be routed to the existing environment hosting the current version of the application.

upvoted 6 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: A

Keywords: minimal disruptions, MINIMAL changes

upvoted 1 times

zijo 9 months ago

Selected Answer: A

I chose A because the question says MINIMAL changes to the application. Canary deployment needs minimal changes to the application as it requires only adding canary settings to the deployment stage. AWS API Gateway supports canary deployments, allowing you to route a percentage of your traffic to a new stage or version of your API.

upvoted 2 times

dkp 11 months ago

Selected Answer: C

C:

Separate Target Group: By introducing a new target group behind the existing Application Load Balancer, you can direct traffic to this new target group without affecting the existing environment. This means the new version of the application can be tested in isolation.

Step-by-Step Traffic Routing: API Gateway allows you to gradually shift user traffic from one target group (existing version) to another (new version). This means you can start with a small percentage of traffic and gradually increase it, allowing you to monitor the new version's performance and stability.

Quick Rollback: If the new version has any issues, you can quickly revert the traffic to the original target group, ensuring minimal disruption to users.

Separate Environment with Canary Deployment: This introduces a completely separate environment, requiring additional infrastructure management and potentially more configuration changes depending on how the environments are set up.

upvoted 2 times

WhyIronMan 11 months, 2 weeks ago

Selected Answer: A

Agree with A. A parallel environment will allow the company to test the deployment, do smoke tests and all the basic stuff to check if the application is working fine. Then, they can start serving traffic to the users, allowing 10% of the users to go to the new environment and test. The key is the word "disruptions" disruption can be considered a failure in the infrastructure, in the communications (networking) but NOT a Bug. Even do, switching 10% of the users to test is best than switching the entire loadbalancer to a new target group because in this scenarios is 100% of users affect against 10%

Details are everything during an investigation...

upvoted 3 times

  **dzn** 1 year ago

Selected Answer: D

A is not meet the "MINIMAL changes to the application" requirement. If application receives requests from a different ALB, the application will receive a different request value, such as HTTP headers, and may need to be modified application. Since D is the same ALB, it is unlikely that changes will be necessary.

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

A is correct: <users experience minimal disruptions during any deployment of a new version of the application.> and <ensure it can quickly roll back updates if there is an issue> means deploy in parallel: canary release or blue/green deployment

B, C and D: If there was a large bug with the a new version, users would experience huge service disruptions

upvoted 3 times

  **Ramdi1** 1 year, 1 month ago

Selected Answer: A

Canary deployment is used to stop disruption hence I have voted A

upvoted 3 times

  **zolthar_z** 1 year, 2 months ago

Selected Answer: D

Answer is D. Even API supports canary deployment it is only if the API can redirect the traffic between two stages, in this case the API sends the traffic directly to the ALB, and from the ALB yo can choose to which environment redirect the traffic

upvoted 2 times

  **Ffida2214** 1 year, 2 months ago

Why not Option D, as the deployment is API gateway-->ALB-->target groups(EC2). and question is saying that: The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application, with minimal changes to application (it is not saying that we shouldn't change deployment steps)



upvoted 1 times

  **OrganizedChaos25** 1 year, 9 months ago

Selected Answer: A

A is definitely correct

upvoted 3 times

  **devnv** 1 year, 9 months ago

A is correct



upvoted 1 times

  **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: A

A is correct answer

upvoted 2 times



  **Sazeka** 1 year, 10 months ago

Selected Answer: A

The correct answer is A

Correct Answer is A. API Gateway supports canary deployment on a deployment stage before you direct all traffic to that stage. A parallel environment means we will create a new ALB and a target group that will target a new set of EC2 instances on which the newer version of the app will be deployed. So the canary setting associated to the new version of the API will connect with the new ALB instance which in turn will direct the traffic to the new EC2 instances on which the newer version of the application is deployed.

upvoted 4 times

  **Jeanphi72** 1 year, 10 months ago

Selected Answer: A

I disagree with C: <https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html>
upvoted 3 times

A company is storing 100 GB of log data in .csv format in an Amazon S3 bucket. SQL developers want to query this data and generate graphs to visualize it. The SQL developers also need an efficient, automated way to store metadata from the .csv file.

Which combination of steps will meet these requirements with the LEAST amount of effort? (Choose three.)

- A. Filter the data through AWS X-Ray to visualize the data.
- B. Filter the data through Amazon QuickSight to visualize the data.
- C. Query the data with Amazon Athena.
- D. Query the data with Amazon Redshift.
- E. Use the AWS Glue Data Catalog as the persistent metadata store.
- F. Use Amazon DynamoDB as the persistent metadata store.

Suggested Answer: BCE

Community vote distribution

BCE (100%)



  **habros**  1 year, 2 months ago

Selected Answer: BCE

BCE. Glue Data Catalog can crawl S3 buckets to store table metadata. Then call the data catalog directly in Athena. It will show the partitions of the data.




Athena does not deal with DynamoDB directly. Hence F is out.

upvoted 6 times

  **habros** 1 year, 2 months ago

<https://docs.aws.amazon.com/glue/latest/dg/catalog-and-crawler.html>

upvoted 3 times

  **thanhv142**  7 months, 1 week ago

BCE are correct: < query this data and generate graphs to visualize it> means athena and quicksight

A: irrelevant

D: too expensive

F: Dynamodb is primarily used for storing web session data and not for this purpose



upvoted 4 times

  **OrganizedChaos25** 1 year, 3 months ago

Selected Answer: BCE


BCE are correct

upvoted 3 times

  **devnv** 1 year, 3 months ago

BCE are correct

upvoted 1 times

  **PhuocT** 1 year, 3 months ago

yep, agree with B,C and E.

upvoted 1 times

  **ParagSanyashiv** 1 year, 3 months ago

Selected Answer: BCE

Agree with BCE

upvoted 2 times

A company deploys its corporate infrastructure on AWS across multiple AWS Regions and Availability Zones. The infrastructure is deployed on Amazon EC2 instances and connects with AWS IoT Greengrass devices. The company deploys additional resources on on-premises servers that are located in the corporate headquarters.

The company wants to reduce the overhead involved in maintaining and updating its resources. The company's DevOps team plans to use AWS Systems Manager to implement automated management and application of patches. The DevOps team confirms that Systems Manager is available in the Regions that the resources are deployed in. Systems Manager also is available in a Region near the corporate headquarters.


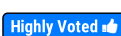
Which combination of steps must the DevOps team take to implement automated patch and configuration management across the company's EC2 instances, IoT devices, and on-premises infrastructure? (Choose three.)

- A. Apply tags to all the EC2 instances, AWS IoT Greengrass devices, and on-premises servers. Use Systems Manager Session Manager to push patches to all the tagged devices.
- B. Use Systems Manager Run Command to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers.
- C. Use Systems Manager Patch Manager to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers as a Systems Manager maintenance window task.
- D. Configure Amazon EventBridge to monitor Systems Manager Patch Manager for updates to patch baselines. Associate Systems Manager Run Command with the event to initiate a patch action for all EC2 instances, AWS IoT Greengrass devices, and on-premises servers.
- E. Create an IAM instance profile for Systems Manager. Attach the instance profile to all the EC2 instances in the AWS account. For the AWS IoT Greengrass devices and on-premises servers, create an IAM service role for Systems Manager.
- F. Generate a managed-instance activation. Use the Activation Code and Activation ID to install Systems Manager Agent (SSM Agent) on each server in the on-premises environment. Update the AWS IoT Greengrass IAM token exchange role. Use the role to deploy SSM Agent on all the IoT devices.

Suggested Answer: BCF

Community vote distribution

CEF (100%)

 **4bed5ff**  1 year, 9 months ago


Selected Answer: CEF

I also choose E instead of B.

Why E is correct: "Previously in this post, you created and deployed the SSM Agent component which would have created an IAM service role. Suppose the AWS IoT Greengrass documentation was followed to deploy the SSM agent. In that case, the name of the IAM service role should be SSMSERVICEROLE."

Why is B wrong: B is redundant given that answer C calls out Systems Manager Patch Manager which itself uses Systems Manager Run Command. Furthermore Run Command is described here to be used to run automated scripts and not to schedule patching: "we'll demonstrate how to use Session Manager to open remote login to an edge device, patch them using Patch Manager, and run automated scripts through Run Command"

Quotes above are from: https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-manager/?force_isolation=true
upvoted 8 times

 **thanhnv142**  1 year, 1 month ago

CEF:

- < implement automated patch > means Systems Manager Patch Manager
- < configuration management > means we need install system manager agent
- we need to configure sufficient permissions for SSM

upvoted 7 times

🗲️ 👤 **jamesf** Most Recent 7 months, 2 weeks ago

Selected Answer: CEF

Systems Manager Patch Manager, System Manager Agent, permission

upvoted 1 times

🗲️ 👤 **c3518fc** 10 months, 3 weeks ago

Selected Answer: CEF

By following the combination of steps C, E, and F, the DevOps team can effectively implement automated patch and configuration management across the company's EC2 instances, IoT Greengrass devices, and on-premises infrastructure using AWS Systems Manager's capabilities and best practices.

upvoted 3 times

🗲️ 👤 **dkp** 11 months ago

Selected Answer: CEF

ans is CEF

upvoted 1 times

🗲️ 👤 **DanShone** 12 months ago

Selected Answer: CEF

CEF are correct

upvoted 1 times

🗲️ 👤 **OrganizedChaos25** 1 year, 9 months ago

Selected Answer: CEF

CEF are correct

upvoted 1 times

🗲️ 👤 **2pk** 1 year, 10 months ago

Agreed with Parag CEF

upvoted 1 times

🗲️ 👤 **ParagSanyashiv** 1 year, 10 months ago

Selected Answer: CEF

CEF make more sense.

upvoted 1 times

🗲️ 👤 **Jeanphi72** 1 year, 10 months ago

Selected Answer: CEF

I disagree with the solution ... FEC for me

upvoted 1 times

A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue/green deployment process with immutable instances when deploying new software.

During testing, users are being automatically logged out of the application at random times. Testers also report that, when a new version of the application is deployed, all users are logged out. The development team needs a solution to ensure users remain logged in across scaling events and application deployments.

What is the MOST operationally efficient way to ensure users remain logged in?

- A. Enable smart sessions on the load balancer and modify the application to check for an existing session.
- B. Enable session sharing on the load balancer and modify the application to read from the session store.
- C. Store user session information in an Amazon S3 bucket and modify the application to read session information from the bucket.
- D. Modify the application to store user session information in an Amazon ElastiCache cluster.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **thanhnv142** Highly Voted 1 year, 1 month ago

D is correct: <During testing, users are being automatically logged out of the application at random times>: the cause is there is no data storage that stores user's session. We need a session data storage to store user session
upvoted 5 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

- A. Enable smart sessions on the load balance: there is no smart session on ALB
 - B. Enable session sharing on the load balancer: load balancer does not store session data
 - C. storing session data in S3 introduces latency
- upvoted 1 times

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

Selected Answer: D

keywords: Amazon ElastiCache cluster
upvoted 1 times

🗳️ 👤 **dkp** 11 months ago

Selected Answer: D

D is correct
upvoted 2 times

🗳️ 👤 **WhyIronMan** 11 months, 2 weeks ago

Selected Answer: D

D is the correct one
upvoted 2 times

🗳️ 👤 **a54b16f** 1 year, 1 month ago

Selected Answer: D

<https://aws.amazon.com/blogs/developer/elasticache-as-an-asp-net-session-store/>
upvoted 1 times

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

Why not C? Compared to D C is serverless thus more operationally efficient.
upvoted 2 times

🗳️ 👤 **dzn** 1 year ago

This is why many web application frameworks support Redis and Memcached session. Also S3 is expensive to read, latency.
upvoted 3 times

🗨️ 👤 **lunt** 1 year, 9 months ago

like comment by accident. S3 is an object store, its not a mounted FS as such, potential performance issues & consistency rule it out. Session data - keep it local, caching system like redis or DB. C actually requires a lot more work as who is now managing the sessions, the bucket, keeping all in sync?

upvoted 2 times

🗨️ 👤 **OrganizedChaos25** 1 year, 9 months ago

Selected Answer: D

D is correct

upvoted 4 times

🗨️ 👤 **devnv** 1 year, 9 months ago

agree with D

upvoted 1 times

🗨️ 👤 **2pk** 1 year, 10 months ago

Yep D, session should be stored to share.

upvoted 1 times

A DevOps engineer needs to configure a blue/green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database. The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.

The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software, blue or green, gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environment's target group or the green environment's target group. An Amazon Route 53 record for www.example.com points to the ALB.

The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environment's EC2 instances.

What should the DevOps engineer do to meet these requirements?

- A. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, use an AWS CLI command to update the ALB to send traffic to the green environment's target group.
- B. Use an AWS CLI command to update the ALB to send traffic to the green environment's target group. Then start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances.
- C. Update the launch template to deploy the green environment's software on the blue environment's EC2 instances. Keep the target groups and Auto Scaling groups unchanged in both environments. Perform a rolling restart of the blue environment's EC2 instances.
- D. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, update the Route 53 DNS to point to the green environment's endpoint on the ALB.

Suggested Answer: A

Community vote distribution

A (100%)

 **PhuocT** Highly Voted 1 year, 9 months ago

A is correct, cannot be D, as there is only one ALB. The ALB can be configured to send traffic to the blue environment's target group or the green environment's target group. Traffic route to blue or green, must be configure at Load balancer, in this case, update the target group.
upvoted 16 times

 **jamesf** Most Recent 7 months, 2 weeks ago

Selected Answer: A

One Application Load Balancer (ALB)
One Auto Scaling Group (ASG) for Blue and one Auto Scaling Group (ASG) for Green
upvoted 2 times

 **dkp** 11 months ago

Selected Answer: A

its A.
upvoted 2 times

 **sirronido** 11 months ago

B correct
option A reverses the order of operations, which goes against the recommended practice of updating the load balancer first to send traffic to the new environment before deploying the new software.
upvoted 1 times

 **thanhnv142** 1 year, 1 month ago

Selected Answer: A

A is correct: <The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environment's EC2 instances.> and <The ALB can be configured to send traffic to the blue environment's target group or the

green environment's target group.>means we should do the traffic migration manually by config the ALB

B and C: no mention of migration step

D: should not use route 53 DNS, we need to configure the ALB

upvoted 4 times

🗳️ 👤 **rif** 1 year, 4 months ago

Answer is A

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html#cfn-attributes-updatepolicy-rollingupdate>

upvoted 2 times

🗳️ 👤 **kyuhobe** 1 year, 6 months ago

The client-side caches the results of DNS queries, so DNS switching lacks immediacy, and it's challenging to transition traffic all at once.

Therefore, option D doesn't meet the requirement of moving traffic all at once and is not suitable.

upvoted 1 times

🗳️ 👤 **ixdb** 1 year, 6 months ago

What???? No one read the question carefully. There are two ALB.

The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment.

upvoted 1 times

🗳️ 👤 **ixdb** 1 year, 6 months ago

D is right.

upvoted 2 times

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

No, you mixed it up. There is only one ALB and two ASGs. Option A is the answer as there is no need to touch Route53 in this scenario(it's pointing to the single ALB already).

upvoted 2 times

🗳️ 👤 **Certified101** 1 year, 7 months ago

Selected Answer: A

Cannot be D as the Route 53 record will be unchanged - points to the same ALB - the target group on the ALB need to be updated

upvoted 2 times

🗳️ 👤 **Blueee** 1 year, 8 months ago

Selected Answer: A

A is correct

upvoted 2 times

🗳️ 👤 **MarDog** 1 year, 8 months ago

I don't think it's D, because it's a single ALB and Route 53 is already pointing at it as www.example.com. What needs to occur is an ASG switch within the ALB. So, A is the best bet.

upvoted 2 times

🗳️ 👤 **walkwolf3** 1 year, 9 months ago

A is correct.

For D, DNS alias record needs to be updated, and green environment's endpoint wasn't mentioned in the question.

upvoted 1 times

🗳️ 👤 **devnv** 1 year, 9 months ago

Yes D sounds the best approach

upvoted 1 times

🗳️ 👤 **2pk** 1 year, 10 months ago

It must be D. DNS record should be diverted to the green ALB

upvoted 1 times

🗳️ 👤 **sb333** 1 year, 7 months ago

The question states there is only one ALB with two target groups (Blue and Green). The DNS record points to that one ALB. So answer D is not correct.

upvoted 1 times

A company is building a new pipeline by using AWS CodePipeline and AWS CodeBuild in a build account. The pipeline consists of two stages. The first stage is a CodeBuild job to build and package an AWS Lambda function. The second stage consists of deployment actions that operate on two different AWS accounts: a development environment account and a production environment account. The deployment stages use the AWS CloudFormation action that CodePipeline invokes to deploy the infrastructure that the Lambda function requires.


A DevOps engineer creates the CodePipeline pipeline and configures the pipeline to encrypt build artifacts by using the AWS Key Management Service (AWS KMS) AWS managed key for Amazon S3 (the aws/s3 key). The artifacts are stored in an S3 bucket. When the pipeline runs, the CloudFormation actions fail with an access denied error.

Which combination of actions must the DevOps engineer perform to resolve this error? (Choose two.)

- A. Create an S3 bucket in each AWS account for the artifacts. Allow the pipeline to write to the S3 buckets. Create a CodePipeline S3 action to copy the artifacts to the S3 bucket in each AWS account. Update the CloudFormation actions to reference the artifacts S3 bucket in the production account.
- B. Create a customer managed KMS key. Configure the KMS key policy to allow the IAM roles used by the CloudFormation action to perform decrypt operations. Modify the pipeline to use the customer managed KMS key to encrypt artifacts.
- C. Create an AWS managed KMS key. Configure the KMS key policy to allow the development account and the production account to perform decrypt operations. Modify the pipeline to use the KMS key to encrypt artifacts.
- D. In the development account and in the production account, create an IAM role for CodePipeline. Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3 bucket. In the CodePipeline account, configure the CodePipeline CloudFormation action to use the roles.
- E. In the development account and in the production account, create an IAM role for CodePipeline. Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3 bucket. In the CodePipeline account, modify the artifacts S3 bucket policy to allow the roles access. Configure the CodePipeline CloudFormation action to use the roles.

Suggested Answer: BD

Community vote distribution



lunt Highly Voted 1 year, 9 months ago

Selected Answer: BE

C = AWS KMS fundamentals. Cannot modify AWS managed KMS key policies. No Cross account access = will not work. Not sure why there is even a discussion on this. Associate level basics.

upvoted 14 times

svjl 1 year, 3 months ago

You can modify the key policies, it is a managed key. What is wrong is change it to use for different account.

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying.html>

upvoted 2 times

robertohyena 1 year, 3 months ago

From your link: <https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-modifying.html>

When changing a key policy, keep in mind the following rules:

- You can view the key policy for an AWS managed key or a customer managed key, but you can only change the key policy for a customer managed key.
- The policies of AWS managed keys are created and managed by the AWS service that created the KMS key in your account.
- You cannot view or change the key policy for an AWS owned key.

upvoted 2 times

heff_bezos 5 months, 2 weeks ago

From your link:

"You can add or remove IAM users, IAM roles, and AWS accounts in the key policy, and change the actions that are allowed or denied for

those principals."

The answer is BE because you don't want to grant permissions to the KMS key for an ENTIRE account, you'd want to allow access for a particular role.

upvoted 1 times

🗳️ 👤 **youonebe** Most Recent 2 months, 2 weeks ago

Selected Answer: BD

there is no need to modify the artifacts S3 bucket policy to allow the roles access

upvoted 1 times

🗳️ 👤 **jamesf** 7 months, 2 weeks ago

Selected Answer: BE

B - Cannot modify AWS managed KMS key policies.

E - Cross account access and we need bucket policies also to be updated, if its same account then we do not need bucket policies permissions

upvoted 2 times

🗳️ 👤 **xdkonorek2** 8 months, 3 weeks ago

Selected Answer: BD

BD,

try it yourself, create account with a bucket, create role with access to s3 operations, and trust policy for another account.

role assumed by another account has full access to s3 resources thereby it's not needed to set up resource policy on s3 bucket

upvoted 3 times

🗳️ 👤 **Venki_dev** 9 months ago

Selected Answer: BD

Answer is BD ,

I have recently implemented similar solution, and my S3 bucket do not have any policy configured , my IAM role has required KMS key permission and it worked.

modifying the S3 bucket policy, but this is not necessary if the IAM roles are correctly configured and used by the CodePipeline CloudFormation action

upvoted 1 times

🗳️ 👤 **Venki_dev** 8 months, 3 weeks ago

I switch to BE ,

because its cross account access and we need bucket policies also to be updated, if its same account then we do not need bucket policies permissions

upvoted 3 times

🗳️ 👤 **c3518fc** 10 months, 3 weeks ago

Selected Answer: BD

Nobody is saying why we are modifying the artifacts in S3 in Option E in the Codecommit account. Doesn't seem to make sense to me.

upvoted 1 times

🗳️ 👤 **dkp** 11 months ago

Selected Answer: BE

BE. are correct

upvoted 2 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

B and E are correct: <fail with an access denied error.> this means there are issues with policies and permissions.

A: no mention of policies

C: This is what the dev team has tried but failed. Can not modify managed key policy, can only view it

D: no mention of configuring S3 bucket policy

upvoted 4 times

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

Selected Answer: BE

exact steps are in this doc

<https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-create-cross-account.html>

upvoted 3 times

🗳️ 👤 **YR4591** 1 year, 4 months ago

Selected Answer: BE

It's BE,

According to this, aws managed kms key can't be used cross account:

<https://repost.aws/knowledge-center/cross-account-access-denied-error-s3>

"Warning: AWS managed AWS KMS key policies can't be modified because they're read-only. However, you can always view both the AWS managed KMS key policies and customer managed KMS key policies. Because AWS managed KMS key policies can't be updated, cross-account permissions also can't be granted for those key policies. Additionally, objects that are encrypted using an AWS managed KMS key can't be accessed by other AWS accounts. For customer managed KMS key policies, you can change the key policy only from the AWS account that created the policy."

upvoted 4 times

🗳️ **Certified101** 1 year, 8 months ago

Selected Answer: BE

BE, bucket policy needs to be amended also as it will assume roles in the prod and dev account

upvoted 3 times

🗳️ **habros** 1 year, 8 months ago

Selected Answer: BE

BE. CMEK = you determine access (key policy) and rotation period (you define instead of 365 days for AWS managed keys). Perfect for cross account resources.

upvoted 2 times

🗳️ **Mail1964** 1 year, 9 months ago

Selected Answer: BE

You can view the key policy for an AWS managed key or a customer managed key, but you can only change the key policy for a customer managed key.

upvoted 4 times

🗳️ **devnv** 1 year, 9 months ago

CE for me

upvoted 1 times

🗳️ **2pk** 1 year, 10 months ago

B & E , i guess too

upvoted 2 times

🗳️ **2pk** 1 year, 9 months ago

I thought again, it should be A & E correct.

B is wrong because The access denied error typically occurs when the IAM roles used by the CloudFormation action lack the necessary permissions to access the required resources. Therefore, option B does not directly address the access denied error in the given scenario.

upvoted 1 times

🗳️ **PhuocT** 1 year, 10 months ago

Selected Answer: BE

I think it is B and E

upvoted 2 times

🗳️ **Jeanphi72** 1 year, 10 months ago

Selected Answer: CE

Questions says: "A DevOps engineer creates the CodePipeline pipeline and configures the pipeline to encrypt build artifacts by using the AWS Key Management Service (AWS KMS) AWS managed key for Amazon S3 (the aws/s3 key)." not CMK ...

upvoted 3 times

🗳️ **sb333** 1 year, 7 months ago

Answer C is incorrect because you cannot "create" an AWS-managed key or modify its key policy. In order to modify a key policy, you need an customer-managed key (Answer B). The question states they used an AWS-managed key, but got an error. So you have to re-evaluate how to make this work, which requires a customer-managed key.

<https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html#key-mgmt>

upvoted 4 times

A company is using an organization in AWS Organizations to manage multiple AWS accounts. The company's development team wants to use AWS Lambda functions to meet resiliency requirements and is rewriting all applications to work with Lambda functions that are deployed in a VPC. The development team is using Amazon Elastic File System (Amazon EFS) as shared storage in Account A in the organization.

The company wants to continue to use Amazon EFS with Lambda. Company policy requires all serverless projects to be deployed in Account B.

A DevOps engineer needs to reconfigure an existing EFS file system to allow Lambda functions to access the data through an existing EFS access point.

Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

- A. Update the EFS file system policy to provide Account B with access to mount and write to the EFS file system in Account A.
- B. Create SCPs to set permission guardrails with fine-grained control for Amazon EFS.
- C. Create a new EFS file system in Account B. Use AWS Database Migration Service (AWS DMS) to keep data from Account A and Account B synchronized.
- D. Update the Lambda execution roles with permission to access the VPC and the EFS file system.
- E. Create a VPC peering connection to connect Account A to Account B.
- F. Configure the Lambda functions in Account B to assume an existing IAM role in Account A.

Suggested Answer: ACE

Community vote distribution



OrganizedChaos25 Highly Voted 1 year, 9 months ago

Selected Answer: ADE

I got ADE

upvoted 12 times

learnwithaniket Highly Voted 1 year, 3 months ago

Selected Answer: ADE

Initially, I thought of A,E,F. But after reading the docs I came to conclusion A,D,E is correct answer.

E: <https://docs.aws.amazon.com/lambda/latest/dg/configuration-filesystem.html#configuration-filesystem-cross-account>

A,D: <https://docs.aws.amazon.com/lambda/latest/dg/configuration-filesystem.html#configuration-filesystem-permissions>

upvoted 7 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: ADE

Should be ADE

VPC peering required.

upvoted 2 times

dkp 11 months ago

Selected Answer: ADE

A,D,E is correct

upvoted 3 times

DanShone 12 months ago

A,D,E is correct

upvoted 3 times

kyuhuck 1 year ago

Selected Answer: AEF

1.need to update the file system plocy on efs to allow mounting the file system into account b

2.need vpc peering between account account a and account b as the pre-requisite

3.need to assume cross-account iam role to describe the mounts so that a specific mount can be chosen

upvoted 1 times

🗳️ 👤 **thanhnv142** 1 year, 1 month ago

Selected Answer: ADE

ADE are correct: <The company wants to continue to use Amazon EFS with Lambda. Company policy requires all serverless projects to be deployed in Account B.> means we need assign relevant IAM policies to lambda in account b

B: no mention of policy

C: no mention of policy

F: <assume an existing IAM role in Account A>: What role?

upvoted 5 times

🗳️ 👤 **a54b16f** 1 year, 1 month ago

Selected Answer: ADE

NOT F: account B will mount EFS and would read/write as a local folder. There is no way/no need to assume role. Option D would assign permission that allow account B to read/write the EFS.

upvoted 5 times

🗳️ 👤 **zain1258** 1 year, 3 months ago

Selected Answer: ADE

It's ADE.

upvoted 3 times

🗳️ 👤 **hzhang** 1 year, 4 months ago

Selected Answer: AEF

D only works if both lamda function and EFS are in the same account.

upvoted 2 times

🗳️ 👤 **zain1258** 1 year, 3 months ago

When peering enabled between two VPCs, this is possible even if the function and EFS are in different account.

upvoted 1 times

🗳️ 👤 **YR4591** 1 year, 4 months ago

Selected Answer: ADE

1) Lambda in account a can get access directly to EFS using cross account policy on the efs.

2) Access to the efs is via network, thats why vpc peering is needed.

<https://aws.amazon.com/blogs/storage/mount-amazon-efs-file-systems-cross-account-from-amazon-eks/>

upvoted 3 times

🗳️ 👤 **RVivek** 1 year, 6 months ago

Selected Answer: AEF

A & E are obvious answers.

D is wrong Lamda execution role is in account B. You cannot directly assign permission to that role . Instead you add AWS STS AssumeRole API call to your Lambda function's code in account B

upvoted 4 times

🗳️ 👤 **sb333** 1 year, 7 months ago

Selected Answer: ADE

<https://docs.aws.amazon.com/efs/latest/ug/create-file-system-policy.html> (Answer A)

<https://aws.amazon.com/blogs/compute/using-amazon-efs-for-aws-lambda-in-your-serverless-applications/> (Answer D)

<https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html> (Answer E)

upvoted 4 times

🗳️ 👤 **unknownuser123** 1 year, 7 months ago

Selected Answer: AEF

AEF Makes more sense

upvoted 3 times

🗳️ 👤 **emupsx1** 1 year, 7 months ago

The answer is AEF because:

A few hours ago, I just finished the DOP-C02 exam.

My score is 1000 points.

This question has come up, I choose AEF.

upvoted 4 times

🗨️ 👤 **CirusD** 1 year, 7 months ago

I am sure you didn't get 1000 if you got this answer wrong

upvoted 3 times

🗨️ 👤 **sb333** 1 year, 7 months ago

Please provide supporting links, since the documentation points to ADE.

<https://docs.aws.amazon.com/efs/latest/ug/create-file-system-policy.html> (Answer A)

<https://aws.amazon.com/blogs/compute/using-amazon-efs-for-aws-lambda-in-your-serverless-applications/> (Answer D)

<https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html> (Answer E)

upvoted 3 times

🗨️ 👤 **sb333** 1 year, 7 months ago

Another support for D and not F.

<https://repost.aws/knowledge-center/access-efs-across-accounts>

This talks about assigning IAM permissions on the account B side, with EFS located in account A. For Lambda, those IAM permissions are part of the execution role. There is nothing indicating the need for using roles from account A. Only an EFS file system policy in account A. And of course peering is needed between the two accounts.

If you did get 1000 points, and you selected AEF, this could have been one of those questions that did not count against your raw score. AWS will have some questions that are not included in your score, but are questions that may be new and are being evaluated.

upvoted 2 times

🗨️ 👤 **zain1258** 1 year, 3 months ago

In exam there are a few questions that does not have any impact on your score. No matter you mark them right or wrong.

upvoted 1 times

🗨️ 👤 **ogwu2000** 1 year, 7 months ago

ADF for me

upvoted 1 times

🗨️ 👤 **ogwu2000** 1 year, 7 months ago

E is wrong . All accounts in same VPC so, you cant do VPC peering.

upvoted 1 times

🗨️ 👤 **devnv** 1 year, 9 months ago

AEF are correct

upvoted 2 times

A media company has several thousand Amazon EC2 instances in an AWS account. The company is using Slack and a shared email inbox for team communications and important updates. A DevOps engineer needs to send all AWS-scheduled EC2 maintenance notifications to the Slack channel and the shared inbox. The solution must include the instances' Name and Owner tags.

Which solution will meet these requirements?

- A. Integrate AWS Trusted Advisor with AWS Config. Configure a custom AWS Config rule to invoke an AWS Lambda function to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe a Slack channel endpoint and the shared inbox to the topic.
- B. Use Amazon EventBridge to monitor for AWS Health events. Configure the maintenance events to target an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe an AWS Lambda function to the SNS topic to send notifications to the Slack channel and the shared inbox.
- C. Create an AWS Lambda function that sends EC2 maintenance notifications to the Slack channel and the shared inbox. Monitor EC2 health events by using Amazon CloudWatch metrics. Configure a CloudWatch alarm that invokes the Lambda function when a maintenance notification is received.
- D. Configure AWS Support integration with AWS CloudTrail. Create a CloudTrail lookup event to invoke an AWS Lambda function to pass EC2 maintenance notifications to Amazon Simple Notification Service (Amazon SNS). Configure Amazon SNS to target the Slack channel and the shared inbox.

Suggested Answer: B

Community vote distribution

B (100%)

dkp 4 months, 4 weeks ago

Selected Answer: B

B is the answer

upvoted 2 times

thanhnv142 7 months, 1 week ago

Selected Answer: B

B is correct: <needs to send all AWS-scheduled EC2 maintenance notifications to the Slack channel and the shared inbox> means SNS

C: no mention of SNS

A: AWS trusted advisor has nothing to do here

D: AWS Support is support plan. It has nothing to do here. so as AWS cloud trail

upvoted 2 times

yuliaqwerty 8 months ago

Answer is B

upvoted 1 times

sarlos 8 months, 1 week ago

Yes it's B

upvoted 1 times

n_d1 1 year, 1 month ago

Selected Answer: B

<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html>

upvoted 4 times

Certified101 1 year, 1 month ago

Selected Answer: B

B is correct



upvoted 2 times

OrganizedChaos25 1 year, 3 months ago

Selected Answer: B

B is the answer I got

upvoted 2 times

  **devnv** 1 year, 3 months ago

B is correct

upvoted 1 times

An AWS CodePipeline pipeline has implemented a code release process. The pipeline is integrated with AWS CodeDeploy to deploy versions of an application to multiple Amazon EC2 instances for each CodePipeline stage.

During a recent deployment, the pipeline failed due to a CodeDeploy issue. The DevOps team wants to improve monitoring and notifications during deployment to decrease resolution times.


What should the DevOps engineer do to create notifications when issues are discovered?

- A. Implement Amazon CloudWatch Logs for CodePipeline and CodeDeploy, create an AWS Config rule to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
- B. Implement Amazon EventBridge for CodePipeline and CodeDeploy, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
- C. Implement AWS CloudTrail to record CodePipeline and CodeDeploy API call information, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
- D. Implement Amazon EventBridge for CodePipeline and CodeDeploy, create an Amazon Inspector assessment target to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

Suggested Answer: B

Community vote distribution

B (100%)

 **thanhnv142** Highly Voted 7 months, 1 week ago

Selected Answer: B

B is correct: <monitoring and notifications during deployment> means eventbridge and SNS

A: cloudwatchlog has nothing to do here. This is use for continuous monitoring of AWS services

C: cloudtrail is for account activities monitoring

D: Inspector is for threat detection


upvoted 5 times

 **dkp** Most Recent 4 months, 4 weeks ago

Selected Answer: B

B is correc

upvoted 2 times

 **YR4591** 10 months, 2 weeks ago

Selected Answer: B

Its B, They want to monitor issued DURING deployment, means near real time, so cloudwatch event will do the work.

C is wrong for to reasons, first, cloudtrail alone can't trigger lambda without an event. Second, cloud trail logs are update in 5 minutes intervals, which means monitoring for the code deploy will not be during deployment.

upvoted 2 times

 **haazybanj** 1 year, 2 months ago

Selected Answer: B

B. Implement Amazon EventBridge for CodePipeline and CodeDeploy, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

Explanation:

Amazon EventBridge provides a serverless event bus that integrates with various AWS services. By implementing EventBridge for CodePipeline and CodeDeploy, the engineer can capture deployment events and trigger actions based on those events. Creating an AWS Lambda function allows for evaluating code deployment issues and performing custom actions. Additionally, creating an Amazon SNS topic provides a means to notify stakeholders of any deployment issues detected.



upvoted 3 times

  **OrganizedChaos25** 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 3 times

  **devnv** 1 year, 3 months ago

Yes it's B

upvoted 1 times

A global company manages multiple AWS accounts by using AWS Control Tower. The company hosts internal applications and public applications.

Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations. One of the AWS Control Tower member accounts serves as a centralized DevOps account with CI/CD pipelines that application teams use to deploy applications to their respective target AWS accounts. An IAM role for deployment exists in the centralized DevOps account.

An application team is attempting to deploy its application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in an application AWS account. An IAM role for deployment exists in the application AWS account. The deployment is through an AWS CodeBuild project that is set up in the centralized DevOps account. The CodeBuild project uses an IAM service role for CodeBuild. The deployment is failing with an Unauthorized error during attempts to connect to the cross-account EKS cluster from CodeBuild.

Which solution will resolve this error?

- A. Configure the application account's deployment IAM role to have a trust relationship with the centralized DevOps account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account's deployment IAM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.
- B. Configure the centralized DevOps account's deployment IAM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the centralized DevOps account's deployment IAM role to allow the required access to CodeBuild.
- C. Configure the centralized DevOps account's deployment IAM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRoleWithSAML action. Configure the centralized DevOps account's deployment IAM role to allow the required access to CodeBuild.
- D. Configure the application account's deployment IAM role to have a trust relationship with the AWS Control Tower management account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account's deployment IAM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Suggested Answer: B

Community vote distribution

A (100%)

 **Certified101** Highly Voted 1 year, 8 months ago

Selected Answer: A

A. Configure the application account's deployment IAM role to have a trust relationship with the centralized DevOps account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account's deployment IAM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Options B, C, and D are not correct because the centralized DevOps account's deployment IAM role doesn't need to trust the application account, it's the other way around. The sts:AssumeRoleWithSAML action in option C is used for federation from a SAML 2.0 compliant identity provider and is not necessary in this scenario. Lastly, there's no need to have a trust relationship with the AWS Control Tower management account as in option D, as the interaction is directly between the DevOps account and the application account.

upvoted 10 times

 **thanhnv142** Highly Voted 1 year, 1 month ago

Selected Answer: A

A is correct: <Unauthorized error during attempts to connect> means we need to setup relevant permissions and policies

- A is correct because < AWS CodeBuild project that is set up in the centralized DevOps account>, so we should setup trust relationship on the account that has resources, which is the application account and allow codebuild from centralized account assume it

B and C are wrong: we need to setup trust from the app account, not the centralized account.

D: this option mentions control Tower, which is irrelevant

upvoted 6 times

 **jamesf** Most Recent 7 months, 1 week ago

Selected Answer: A

A. Configure the application account's deployment IAM role to have a trust relationship with the centralized DevOps account.

- setup trust relationship on the account that has resources, which is the application account

Configure the trust relationship to allow the sts:AssumeRole action.

- allow CodeBuild from centralized account assume it

- CodeBuild is configured in Centralized DevOps account but not in application account.

Configure the application account's deployment IAM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

- the application account has access to the resources

upvoted 2 times

  **tartarus23** 1 year, 8 months ago

Selected Answer: A

(A) This solution addresses the Unauthorized error by allowing the DevOps account to assume the IAM role in the application account that has the necessary permissions to access the EKS cluster. The other options don't provide the necessary cross-account permissions or correctly configure the roles for accessing EKS.

upvoted 3 times

  **walkwolf3** 1 year, 9 months ago

B is correct.

Unauthorized error happened from CodeBuild in Dev account to EKS cluster in application account, instead of reverse direction.

upvoted 2 times

  **zain1258** 1 year, 3 months ago



CodeBuild is configured in Centralized DevOps account not in application account.

upvoted 2 times

  **2pk** 1 year, 9 months ago

I'd like to add more, don't get the source and destination mixed up. Because the Application team must deploy resources in the Dev account. So, the source should be the Application team and the destination should be the Dev team.

upvoted 2 times

  **PhuocT** 1 year, 9 months ago

Selected Answer: A



A is correct

upvoted 1 times

  **ParagSanyashiv** 1 year, 10 months ago

A is correct Answer

upvoted 1 times

  **2pk** 1 year, 10 months ago

Answer is A.

In the source AWS account, the IAM role used by the CI/CD pipeline should have permissions to access the source code repository, build artifacts, and any other resources required for the build process.

In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application.

the IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

upvoted 3 times


A highly regulated company has a policy that DevOps engineers should not log in to their Amazon EC2 instances except in emergencies. If a DevOps engineer does log in, the security team must be notified within 15 minutes of the occurrence.

Which solution will meet these requirements?

- A. Install the Amazon Inspector agent on each EC2 instance. Subscribe to Amazon EventBridge notifications. Invoke an AWS Lambda function to check if a message is about user logins. If it is, send a notification to the security team using Amazon SNS.
- B. Install the Amazon CloudWatch agent on each EC2 instance. Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user logins. If a login is found, send a notification to the security team using Amazon SNS.
- C. Set up AWS CloudTrail with Amazon CloudWatch Logs. Subscribe CloudWatch Logs to Amazon Kinesis. Attach AWS Lambda to Kinesis to parse and determine if a log contains a user login. If it does, send a notification to the security team using Amazon SNS.
- D. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3. Set up an S3 event to invoke an AWS Lambda function, which invokes an Amazon Athena query to run. The Athena query checks for logins and sends the output to the security team using Amazon SNS.

Suggested Answer: B

Community vote distribution



thanhv142 Highly Voted 1 year, 1 month ago

Selected Answer: B

B is correct: <the security team must be notified > means SNS

A: irrelevant, inspector is for vulnerability scanning

C: cloud trail is for monitoring account activities

D: This options uses manual script, which is irrelevant

upvoted 6 times

zijo Most Recent 8 months, 2 weeks ago

Selected Answer: B

B is the cheapest and correct solution

CloudTrail does not capture:

SSH logins to Linux instances.

RDP logins to Windows instances.

Commands executed on the instances.

Local file access or modifications within the instances.

upvoted 1 times

haazybanj 1 year, 7 months ago

Selected Answer: C

While Option B can provide valuable insights into user logins and send notifications to the security team, it might not guarantee that the security team is notified within 15 minutes of a login occurrence. The time it takes for the CloudWatch metric filter to process and detect the login event, along with the potential delays in the SNS notification, could result in notifications being sent beyond the required 15-minute timeframe.

On the other hand, Option C, which uses AWS CloudTrail with Amazon CloudWatch Logs and Amazon Kinesis, allows real-time processing and immediate notifications when a user login event is detected. This makes Option C more suitable for meeting the specific requirement of notifying the security team within 15 minutes of a login occurrence.

upvoted 1 times

RVivek 1 year, 5 months ago

Cloud Trail will track calls to AWS API, but not the OS login in an EC2. That can be checked only using Cloud watch logs

upvoted 10 times

n_d1 1 year, 7 months ago

Selected Answer: B

<https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-instances/>

upvoted 3 times

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

Selected Answer: B

B,

Eventhough it's not stated in some questions, the cheapest solution to a problem is always AWS favorite.

upvoted 4 times

🗲️ 👤 **gdtypk** 1 year, 9 months ago

Selected Answer: B

Isn't it possible to get login events with CloudTrail?

upvoted 2 times

🗲️ 👤 **Mail1964** 1 year, 9 months ago

Selected Answer: B

Subtle difference Cloudtrail is "near" realtime - You can use subscriptions to get access to a real-time feed of log events from CloudWatch Logs and have it delivered to other services such as an Amazon Kinesis stream, an Amazon Kinesis Data Firehose stream, or AWS Lambda for custom processing, analysis, or loading to other systems.

upvoted 3 times

🗲️ 👤 **devnv** 1 year, 9 months ago

B is the right answer

upvoted 1 times

🗲️ 👤 **2pk** 1 year, 10 months ago

i think its C, Both B&C solutions are valid and can meet the requirement of notifying the security team within 15 minutes of a DevOps engineer logging into an EC2 instance.

However, there are some differences in how quickly each solution can detect and notify the security team of a login event.

The CloudTrail-based solution can detect a login event more quickly than the CloudWatch-based solution because CloudTrail captures API events in near-real-time, while CloudWatch logs may have a delay of a few minutes before they appear in the log group. Therefore, the CloudTrail-based solution is more likely to meet the 15-minute notification requirement.

upvoted 1 times

🗲️ 👤 **buiquangbk90** 1 year, 6 months ago

AWS CloudTrail captures API calls made on your account and sends log files to CloudWatch Logs. The provided solution monitors for login-related API calls. While this may detect some login activity (like a RunInstances API call), it will not catch SSH logins to an EC2 instance. Therefore, this solution isn't comprehensive enough.

=> Correct answer is B.

upvoted 3 times

A company updated the AWS CloudFormation template for a critical business application. The stack update process failed due to an error in the updated template, and AWS CloudFormation automatically began the stack rollback process. Later, a DevOps engineer discovered that the application was still unavailable and that the stack was in the UPDATE_ROLLBACK_FAILED state.

Which combination of actions should the DevOps engineer perform so that the stack rollback can complete successfully? (Choose two.)

- A. Attach the AWSCloudFormationFullAccess IAM policy to the AWS CloudFormation role.
- B. Automatically recover the stack resources by using AWS CloudFormation drift detection.
- C. Issue a ContinueUpdateRollback command from the AWS CloudFormation console or the AWS CLI.
- D. Manually adjust the resources to match the expectations of the stack.
- E. Update the existing AWS CloudFormation stack by using the original template.

Suggested Answer: CD

Community vote distribution

CD (100%)

  **2pk** Highly Voted 1 year, 10 months ago

yes C & D

C. Issue a ContinueUpdateRollback command from the AWS CloudFormation console or the AWS CLI: This command allows CloudFormation to continue the rollback process from the point where it was paused. By using this command, CloudFormation will attempt to restore the resources to their previous state and delete any resources that were created during the update.

D. Manually adjust the resources to match the expectations of the stack: This involves identifying and correcting the root cause of the update failure, which could involve changing the resource configuration or resolving any dependencies or inconsistencies in the stack.

upvoted 10 times

  **heff_bezos** Most Recent 5 months, 2 weeks ago

Selected Answer: CD

D.

"In most cases, you must fix the error that causes the update rollback to fail before you can continue to roll back your stack. In other cases, you can continue to roll back the update without any changes, for example when a stack operation times out."

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-continueupdaterollback.html>

upvoted 2 times

  **thanhnv142** 1 year, 1 month ago

C and D are correct: < UPDATE_ROLLBACK_FAILED state> means we are left with ContinueUpdateRollback command

A: irrelevant, AWSCloudFormationFullAccess IAM policy is used to access ACF, not to fix this

B: AWS CloudFormation drift detection is to check if the stack has been updated unexpectedly, not to fix irrelevant

E: The original template didnt work, so update the stack using it wont work



upvoted 3 times

  **yuliaqwerty** 1 year, 2 months ago

Selected Answer: CD



Agree with C and D

upvoted 2 times

  **sarlos** 1 year, 2 months ago

C and D are right

upvoted 1 times

  **Certified101** 1 year, 7 months ago

Selected Answer: CD

Yes it's C & D

upvoted 3 times



  **tartarus23** 1 year, 8 months ago

Selected Answer: CD

(C) This command will try to roll back the stack to the previously working state after you have resolved the issues that caused the rollback failure.

(D) Sometimes a stack update can fail because the current state of a resource differs from the state expected by AWS CloudFormation (e.g., a resource that AWS CloudFormation is trying to modify or delete is locked by another process). Manually resolving these issues, by stopping the conflicting process or by modifying the resource to match the expected state, will allow the stack update or rollback to proceed.

upvoted 4 times

  **devnv** 1 year, 9 months ago

Yes it's C & D

upvoted 2 times

A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this, downloads the artifact from Amazon S3, and unzips the artifact to complete the deployment.

A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.


Which combination of actions will accomplish this? (Choose three.)

- A. Allow developers to check the code into a code repository. Using Amazon EventBridge, on every pull into the main branch, invoke an AWS Lambda function to build the artifact and store it in Amazon S3.
- B. Create a custom script to clear the cache. Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.
- C. Create user data for each Amazon EC2 instance that contains the clear cache script. Once deployed, test the application. If it is not successful, deploy it again.
- D. Set up AWS CodePipeline to deploy the application. Allow developers to check the code into a code repository as a source for the pipeline.
- E. Use AWS CodeBuild to build the artifact and place it in Amazon S3. Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.
- F. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

Suggested Answer: BDE

Community vote distribution

BDE (100%)

 **tartarus23** Highly Voted 1 year, 2 months ago


Selected Answer: BDE

(B) This would help ensure that the local cache is cleared before the new version of the application is installed. AppSpec (Application Specification) file is a unique file to AWS CodeDeploy. It defines the deployment actions you want AWS CodeDeploy to execute.

(D) This would allow you to automate the build and deployment processes. AWS CodePipeline is a fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates.

(E) This would allow you to automate the build process and ensure that the application is built in a consistent environment. AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy. AWS CodeDeploy automates software deployments to a variety of compute services including Amazon EC2, AWS Fargate, AWS Lambda, and your on-premises servers.


upvoted 8 times

 **Certified101** Highly Voted 1 year, 1 month ago

Selected Answer: BDE

BDE combination will do all requirements

upvoted 5 times

 **thanhnv142** Most Recent 7 months, 1 week ago

Selected Answer: BDE

BDE are correct: < migrate to a CI/CD process > means codepipeline, code build and code deploy


A,C and F: no mention of the above three

upvoted 3 times

 **sarlos** 8 months, 1 week ago

BDE is right

upvoted 1 times

 **devnv** 1 year, 3 months ago

Yes it's BD&E

upvoted 2 times

A DevOps engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps manager has been asked to review the company buildspec.yaml file for an AWS CodeBuild project and provide recommendations. The buildspec.yaml file is configured as follows:

```
env:
  variables:
    AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
    AWS_SECRET_ACCESS_KEY: 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
    AWS_DEFAULT_REGION: us-east-1
    DB_PASSWORD: cuj5RptFa3va
  phases:
    build:
      commands:
        - aws s3 cp s3://db-deploy-bucket/my.cnf.template/tmp/my.cnf
        - sed -i 's/DB_PW/${DB_PASSWORD}/' /tmp/my.cnf
        - aws s3 cp s3://db-deploy-bucket/instance.key /tmp/instance.key
        - chmod 600 /tmp/instance.key
        - scp -i /tmp/instance.key /tmp/my.cnf root@10.25.15.23:/etc/my.cnf
        - ssh -i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

What changes should be recommended to comply with AWS security best practices? (Choose three.)

- A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.
- B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.
- C. Store the DB_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables.
- D. Move the environment variables to the 'db-deploy-bucket' Amazon S3 bucket, add a prebuild stage to download, then export the variables.
- E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.
- F. Scramble the environment variables using XOR followed by Base64, add a section to install, and then run XOR and Base64 to the build phase.

Suggested Answer: BCE

Community vote distribution

BCE (77%)

ABC (23%)

WhyIronMan Highly Voted 11 months, 2 weeks ago

Selected Answer: BCE

BCE is correct

A is WRONG. CodeBuild does not keep files for next builds in that way, once the build is done, the files will be deleted.

upvoted 10 times

sb333 Highly Voted 1 year, 7 months ago

Selected Answer: BCE

BCE are the correct answers.

upvoted 5 times

heff_bezos Most Recent 5 months, 2 weeks ago

Selected Answer: BCE

Code Build is a managed service. There's no way for other users to see what's in the container.

upvoted 2 times

jamesf 7 months, 1 week ago

Selected Answer: BCE

Prefer BCE

Option A incorrect as

- CodeBuild does not keep files for next builds in that way, once the build is done, the files will be deleted.
- and don't think have such "CodeBuild users"

upvoted 1 times

🗄️ 👤 **ericphl** 7 months, 3 weeks ago

Selected Answer: ABC

ABC seems right.

upvoted 1 times

🗄️ 👤 **ajeeshb** 8 months, 1 week ago

Selected Answer: ABC

A - Cleans up temp files that stores the my.cnf and the instance key files

B - Removes hardcoded AWS credentials

C - Securely stores DB password

upvoted 1 times

🗄️ 👤 **Diego1414** 1 year ago

Selected Answer: ABC

ABC seems appropriate, since the emphasis is on security.

upvoted 2 times

🗄️ 👤 **thanhnv142** 1 year, 1 month ago

Selected Answer: ABC

ABC are correct: security best practices are related to removing credentials and sensitive data

- A remove temporary files is important because they might contain sensitive data

- B: <remove the AWS credentials> is removing the access key

- C: <remove the DB_PASSWORD> means removing hardcoded DB_PASSWORD

All other options dont relate to sensitive data or password

upvoted 2 times

🗄️ 👤 **sarlos** 1 year, 2 months ago

its BCE

<https://stackoverflow.com/questions/76854227/i-want-to-copy-files-to-aws-ec2-using-buildspec-yml-file-the-22-port-is-open-fo>

upvoted 1 times

🗄️ 👤 **zain1258** 1 year, 3 months ago

Selected Answer: BCE

It's BCE.

A is wrong. I don't think there is any concept of `CodeBuild users`.

upvoted 4 times

🗄️ 👤 **buiquangbk90** 1 year, 6 months ago

BCE

<https://www.examttopics.com/discussions/amazon/view/46729-exam-aws-devops-engineer-professional-topic-1-question-17/>

upvoted 1 times

🗄️ 👤 **einn** 1 year, 7 months ago

Selected Answer: ABC

A: remove sensitive data that could left behind in container

B: remove credentials and use role

C: Use SecureString AWS Systems Manager Parameter Store

upvoted 1 times

🗄️ 👤 **Certified101** 1 year, 7 months ago

Selected Answer: BCE

BCE are the correct ones.

upvoted 3 times

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

Selected Answer: BCE



BCE are the correct ones.

upvoted 3 times

🗄️ 👤 **Kodoma** 1 year, 9 months ago



BCE is correct answer

upvoted 1 times

  **devnv** 1 year, 9 months ago

Sorry, I've read again and it's AB & C.

upvoted 3 times

  **devnv** 1 year, 9 months ago

Yes BCE are correct

upvoted 1 times

A company has a legacy application. A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference.


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

- A. Create a custom Docker image that contains all the dependencies for the legacy application. Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository. Configure a new AWS CodeBuild project to use the custom Docker image to build the deployable artifact and to save the artifact to the S3 bucket.
- B. Launch a new Amazon EC2 instance. Install all the dependencies for the legacy application on the EC2 instance. Use the EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
- C. Create a custom EC2 Image Builder image. Install all the dependencies for the legacy application on the image. Launch a new Amazon EC2 instance from the image. Use the new EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
- D. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with an AWS Fargate profile that runs in multiple Availability Zones. Create a custom Docker image that contains all the dependencies for the legacy application. Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository. Use the custom Docker image inside the EKS cluster to build the deployable artifact and to save the artifact to the S3 bucket.

Suggested Answer: A

Community vote distribution

A (100%)

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

Selected Answer: A

The most operationally efficient solution for automating the process of building the deployable artifact for the legacy application and storing it in an existing Amazon S3 bucket is:

A. This solution leverages containerization with Docker, which allows for consistent and isolated builds, making it easier to manage application dependencies. The use of AWS CodeBuild allows for scalable and automated builds using the custom Docker image, making the process efficient and reliable. The deployable artifact can then be saved to the existing S3 bucket for future reference and deployments.

upvoted 8 times

 **thanhnv142** Highly Voted 1 year, 1 month ago

Selected Answer: A

A is correct: <needs to automate the process of building the deployable artifact for the legacy application> means codebuild
BCD dont mention codebuild, only A mentions

upvoted 6 times

 **jamesf** Most Recent 7 months, 1 week ago

Selected Answer: A

keywords: CodeBuild for Reusable artifacts


upvoted 1 times

 **habros** 1 year, 8 months ago

Selected Answer: A

Reusable artifacts = A.

upvoted 3 times

 **FunkyFresco** 1 year, 8 months ago

Selected Answer: A

Option A makes more sense to me.

upvoted 2 times

 **tartarus23** 1 year, 8 months ago

(A) This approach is the most operationally efficient because it leverages the benefits of containerization, such as isolation and reproducibility, as well as AWS managed services. AWS CodeBuild is a fully managed build service that can compile your source code, run tests, and produce deployable software packages. By using a custom Docker image that includes all dependencies, you can ensure that the environment in which

your code is built is consistent. Using Amazon ECR to store Docker images lets you easily deploy the images to any environment. Also, you can directly upload the build artifacts to Amazon S3 from AWS CodeBuild, which is beneficial for version control and archival purposes.

upvoted 3 times

A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an IAM service role that has permissions to access the S3 bucket.

A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the IAM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.

When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.


Which solution will resolve the issue of failed access to the ECR repository?

- A. Update the buildspec.yml file to log in to the ECR repository by using the `aws ecr get-login-password` AWS CLI command to obtain an authentication token. Update the docker login command to use the authentication token to access the ECR repository.
- B. Add an environment variable of type `SECRETS_MANAGER` to the CodeBuild project. In the environment variable, include the ARN of the CodeBuild project's IAM service role. Update the buildspec.yml file to use the new environment variable to log in with the docker login command to access the ECR repository.
- C. Update the ECR repository to be a public image repository. Add an ECR repository policy that allows the IAM service role to have access.
- D. Update the buildspec.yml file to use the AWS CLI to assume the IAM service role for ECR operations. Add an ECR repository policy that allows the IAM service role to have access.

Suggested Answer: A

Community vote distribution


A (100%)

 **tartarus23** Highly Voted 1 year, 2 months ago

Selected Answer: A

(A) When Docker communicates with an Amazon Elastic Container Registry (ECR) repository, it requires authentication. You can authenticate your Docker client to the Amazon ECR registry with the help of the AWS CLI (Command Line Interface). Specifically, you can use the "aws ecr get-login-password" command to get an authorization token and then use Docker's "docker login" command with that token to authenticate to the registry. You would need to perform these steps in your buildspec.yml file before attempting to push or pull images from/to the ECR repository.

upvoted 7 times

 **haazybanj** Highly Voted 1 year, 1 month ago

Selected Answer: A

A:

When using Amazon ECR, you need to authenticate Docker to the ECR registry before pushing or pulling container images. The authentication token can be obtained using the `aws ecr get-login-password` AWS CLI command. The obtained token needs to be used with the `docker login` command to authenticate Docker to the ECR repository.

By following this approach, the CodeBuild project will have the necessary credentials to access the ECR repository, and the build job will be able to push the container image to the ECR repository successfully.

upvoted 5 times

 **thanhnv142** Most Recent 7 months, 1 week ago

Selected Answer: A

A is correct: <the job fails when the job tries to access the ECR repository.> This means there is problem when accessing the repo. <adjusts the IAM service role with the permissions that are necessary to work with the new ECR repository> means have got sufficient permission. Need token to access with `aws ecr get-login-password` command

BCD no mention of `ecr get-login-password`



upvoted 4 times

 **ixdb** 1 year ago

Selected Answer: A

A is right.

upvoted 3 times

  **CirusD** 1 year, 1 month ago

A..version: 0.2

phases:

pre_build:

commands:

- \$(aws ecr get-login --no-include-email --region region-name)

build:

commands:

- docker build -t repository-name .

- docker tag repository-name:latest repository-uri:latest

post_build:

commands:

- docker push repository-uri:latest

upvoted 3 times

A company manually provisions IAM access for its employees. The company wants to replace the manual process with an automated process. The company has an existing Active Directory system configured with an external SAML 2.0 identity provider (IdP).


The company wants employees to use their existing corporate credentials to access AWS. The groups from the existing Active Directory system must be available for permission management in AWS Identity and Access Management (IAM). A DevOps engineer has completed the initial configuration of AWS IAM Identity Center (AWS Single Sign-On) in the company's AWS account.

What should the DevOps engineer do next to meet the requirements?

- A. Configure an external IdP as an identity source. Configure automatic provisioning of users and groups by using the SCIM protocol.
- B. Configure AWS Directory Service as an identity source. Configure automatic provisioning of users and groups by using the SAML protocol.
- C. Configure an AD Connector as an identity source. Configure automatic provisioning of users and groups by using the SCIM protocol.
- D. Configure an external IdP as an identity source. Configure automatic provisioning of users and groups by using the SAML protocol.

Suggested Answer: A

Community vote distribution



tartarus23 Highly Voted 1 year, 8 months ago

Selected Answer: A

(A) AWS SSO (Single Sign-On) integrates with external identity providers using SAML 2.0, and it can automatically synchronize users and groups from a connected directory using the SCIM (System for Cross-domain Identity Management) protocol. Thus, the DevOps engineer should configure the external IdP as an identity source and then configure automatic provisioning of users and groups by using the SCIM protocol. This will ensure the groups from the existing Active Directory system are available for permission management in AWS Identity and Access Management (IAM) and that employees can use their existing corporate credentials to access AWS.

upvoted 8 times

jamesf Most Recent 7 months, 1 week ago

Selected Answer: A

For Note: SAML (Security Assertion Markup Language) is primarily used for authentication and authorization while SCIM (System for Cross-domain Identity Management) is a protocol used for automating user provisioning and deprovisioning across different systems and domains

upvoted 3 times

thanhnv142 1 year, 1 month ago

Selected Answer: A

A is correct: <The company wants employees to use their existing corporate credentials to access AWS> means we need to assign the existing IdP as an identity source

B: <Configure AWS Directory Service as an identity source> is irrelevant

C: <Configure an AD Connector as an identity source>: AD connector is used for connecting AWS active directory with that of on-prem. This question requires AWS identity Center

D: <provisioning of users and groups by using the SAML protocol>: SAML is an authentication protocol. SCIM is the protocol for IdP connection

upvoted 4 times

zolthar_z 1 year, 3 months ago

Selected Answer: A

A: Explanation: What is the difference between SCIM and SSO? SSO (single-sign on) is a way to authenticate (sign in), and SCIM is a way to provision (create an account).

upvoted 2 times

XP_2600 1 year, 6 months ago

This is quoted from AWS documentation: The SAML protocol however does not provide a way to query the IdP to learn about users and groups. Therefore, you must make IAM Identity Center aware of those users and groups by provisioning them into IAM Identity Center.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/scim-profile-saml.html>

upvoted 1 times

🗲️ 👤 **CirusD** 1 year, 7 months ago

Answer is A : AWS Single Sign-On (AWS SSO) can be integrated with an external SAML 2.0 identity provider (IdP). AWS SSO also supports automatic provisioning (auto-provisioning) of user and group information using the System for Cross-domain Identity Management (SCIM) protocol.

upvoted 1 times

🗲️ 👤 **sb333** 1 year, 7 months ago

Selected Answer: A

Answer A is correct. It is SCIM that can provision users and groups in AWS. Of course the IdP needs to support SCIM (AWS has a list of IdPs that use SCIM). Answer D is not correct as SAML is an authentication protocol (cannot be used to provision users in AWS).

<https://docs.aws.amazon.com/singlesignon/latest/userguide/scim-profile-saml.html>

<https://docs.aws.amazon.com/singlesignon/latest/userguide/supported-idps.html>

upvoted 2 times

🗲️ 👤 **haazybanj** 1 year, 7 months ago

Selected Answer: A

The AWS IAM Identity Center (AWS Single Sign-On) has been configured initially. Now, to automate the provisioning of users and groups from the external IdP into AWS IAM, the engineer should choose the SCIM protocol. SCIM is specifically designed for automatic user provisioning, making it the appropriate choice for this scenario.

Option D (Configure an external IdP as an identity source and use the SAML protocol) could work, but it does not address the requirement for automatic provisioning of users and groups. The use of SCIM (Option A) is preferred for automated user and group provisioning, as it is designed for this purpose.

upvoted 1 times

🗲️ 👤 **Snape** 1 year, 7 months ago

Selected Answer: D

The company already has an external SAML 2.0 IdP, so the DevOps engineer should configure this IdP as an identity source in AWS Single Sign-On. Vs in option A would require to configure new identity source

upvoted 1 times

🗲️ 👤 **habros** 1 year, 8 months ago

Selected Answer: A

A. SCIM is the automated way to provision users. You do it in AAD/AD and it propagates automatically into AWS SSO.

upvoted 1 times

🗲️ 👤 **Blueee** 1 year, 8 months ago

Selected Answer: A

SCIM protocol is to sync the user and groups from the external identity source

upvoted 2 times

🗲️ 👤 **Toptip** 1 year, 8 months ago

Selected Answer: D

D is correct

upvoted 1 times

A company is using AWS to run digital workloads. Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.

The company wants to enforce security standards across the entire organization. To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation. A production support team can modify resources in the production environment by using the AWS Management Console to troubleshoot and resolve application-related issues.


A DevOps engineer must implement a solution to identify in near real time any AWS service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.

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

- A. Use CloudFormation drift detection to identify noncompliant resources. Use drift detection events from CloudFormation to invoke an AWS Lambda function for remediation. Configure the Lambda function to publish logs to an Amazon CloudWatch Logs log group. Configure an Amazon CloudWatch dashboard to use the log group for tracking.
- B. Turn on AWS CloudTrail in the AWS accounts. Analyze CloudTrail logs by using Amazon Athena to identify noncompliant resources. Use AWS Step Functions to track query results on Athena for drift detection and to invoke an AWS Lambda function for remediation. For tracking, set up an Amazon QuickSight dashboard that uses Athena as the data source.
- C. Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resources. Enable AWS Security Hub with the --no-enable-default-standards option in all the AWS accounts. Set up AWS Config managed rules and custom rules. Set up automatic remediation by using AWS Config conformance packs. For tracking, set up a dashboard on Security Hub in a designated Security Hub administrator account.
- D. Turn on AWS CloudTrail in the AWS accounts. Analyze CloudTrail logs by using Amazon CloudWatch Logs to identify noncompliant resources. Use CloudWatch Logs filters for drift detection. Use Amazon EventBridge to invoke the Lambda function for remediation. Stream filtered CloudWatch logs to Amazon OpenSearch Service. Set up a dashboard on OpenSearch Service for tracking.

Suggested Answer: C

Community vote distribution



heff_bezos 5 months, 2 weeks ago

Selected Answer: C

I don't think it is A because the question is asking the LEAST development overhead. Configuring Lambdas to remediate and send logs is development. It is much easier to use the built in features of AWS Config and SecurityHub
upvoted 1 times

heff_bezos 5 months, 2 weeks ago

Lambda functions also have a execution limit of 15 minutes. If a remediation task were to take longer than that, it would fail.
upvoted 1 times

zijo 8 months, 1 week ago

Selected Answer: C

C is the better solution. AWS CloudFormation drift detection helps identify whether the actual configuration of your AWS resources matches their expected configuration as defined in the CloudFormation stack template. While it is a powerful tool for maintaining compliance and consistency, it alone cannot fully prevent noncompliance due to security misconfigurations. That's where you need AWS Config to continuously monitor service configurations and even use aggregator to collect all aws config data from all member accounts in aws organization to Security Hub to provide a centralized dashboard.
upvoted 3 times

Gomer 8 months, 3 weeks ago

Leaning towards "A" unless someone can convince me otherwise. Why?:

I have a problem with this step in "C": "Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resources."

The fact is your not going to detect any "drift" by turning on the recorder AFTER the accounts are noncompliant.



AWS Config rules (canned or custom) and Conformance Packs can do a lot, but it's definitely duplicating settings any security settings already defined CloudFormation stacks.

I lean towards "A" because "To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation".

Therefore CloudFormation stacks are is where the security settings are defined, and thereby CloudFormation is implied to be part of the detection and remediation process.

CloudFormation drift detection can be automated, and one can just "automatically remediate the issue within 15 minutes of identification" by just doing a stack refresh. Easy peasy.

upvoted 2 times

  **ajeeshb** 8 months, 1 week ago

Correct, A is the answer.

upvoted 1 times

  **dkp** 11 months ago

Selected Answer: C

answer is C with minimal overhead

upvoted 1 times

  **tristan_07** 11 months, 4 weeks ago

Selected Answer: A

Both Option A and C work. However, considering Option C involves a lot 'all the AWS accounts,' it undoubtedly increases development overhead

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago


Selected Answer: A

A is correct: drift detection is the best for this scenario, which utilizes AWS cloudformation

B and D: using cloudtraid is for monitoring account activities

C: AWS Config conformance packs cannot make remediation actions. It needs to trigger AWS SSM automation document

upvoted 1 times

  **vn_thanhtung** 9 months, 3 weeks ago

A not correct because not mention how to remediate



upvoted 1 times

  **AzureDP900** 1 year, 5 months ago

C is right

<https://aws.amazon.com/blogs/security/optimize-aws-config-for-aws-security-hub-to-effectively-manage-your-cloud-security-posture/>

upvoted 2 times

  **Snape** 1 year, 7 months ago

Selected Answer: C

Compliance usually indicates towards config

upvoted 3 times

  **ds50421** 1 year, 8 months ago

Selected Answer: C

compliance means aws config,automatic remedy aws config,central dashboard security hub

upvoted 2 times

  **tartarus23** 1 year, 8 months ago

Selected Answer: C

(C) This solution meets all of the requirements. AWS Config can monitor resource configurations for compliance with defined rules. The use of AWS Security Hub allows for centralized management of security alerts and compliance checks across all accounts. AWS Config conformance packs allow for automated remediation of non-compliant resources. AWS Security Hub provides a comprehensive view of high-priority security alerts and compliance status across AWS accounts. This solution is also the one with the least development overhead as it uses built-in AWS services specifically designed for configuration management and compliance tracking.

upvoted 4 times

A company uses AWS Organizations to manage its AWS accounts. The organization root has an OU that is named Environments. The Environments OU has two child OUs that are named Development and Production, respectively.

The Environments OU and the child OUs have the default FullAWSAccess policy in place. A DevOps engineer plans to remove the FullAWSAccess policy from the Development OU and replace the policy with a policy that allows all actions on Amazon EC2 resources.

What will be the outcome of this policy replacement?

- A. All users in the Development OU will be allowed all API actions on all resources.
- B. All users in the Development OU will be allowed all API actions on EC2 resources. All other API actions will be denied.
- C. All users in the Development OU will be denied all API actions on all resources.
- D. All users in the Development OU will be denied all API actions on EC2 resources. All other API actions will be allowed.

Suggested Answer: B

Community vote distribution

B (78%)

A (22%)

  **d262e67** Highly Voted 1 year, 2 months ago


Selected Answer: B

The key point is that "SCP inheritance works differently for Allow and Deny policies". Allowed policies are only inherited if the children don't have any Allow policy. Once they have an allow policy, only actions defined in that policy will be allowed and no "Allow" policy will be inherited from the parent(s) OUs. What inherits is the implicit Deny policy which is a hidden policy sitting above all.

Check the tables in this link:



<https://aws.amazon.com/blogs/security/get-more-out-of-service-control-policies-in-a-multi-account-environment/>

upvoted 13 times

  **MalonJay** 10 months, 1 week ago

Very good link about SCPs.

upvoted 1 times

  **devakram** Highly Voted 11 months ago

Selected Answer: B

I've just tested in my AWS account with the same scenario. I removed the SCP from the dev env and kept the EC2 policy, which by that I was denied access to all other operations except EC2.

upvoted 7 times

  **auxwww** Most Recent 4 months, 4 weeks ago

Selected Answer: B

Best explanation I found in this forum

From: learnwithaniket

"For a permission to be allowed for a specific account, there must be an explicit Allow statement at every level from the root through each OU in the direct path to the account (including the target account itself). This is why when you enable SCPs, AWS Organizations attaches an AWS managed SCP policy named FullAWSAccess which allows all services and actions. If this policy is removed and not replaced at any level of the organization, all OUs and accounts under that level would be blocked from taking any actions.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html

upvoted 1 times

  **HayLLiHuK** 11 months, 1 week ago

Selected Answer: B

Note: Adding an SCP with full AWS access doesn't give all the principals in an account access to everything. SCPs don't grant permissions; they are used to filter permissions. Principals still need a policy within the account that grants them access.

upvoted 2 times

  **DanShone** 12 months ago

Selected Answer: A


A - Inherited SCPs cannot be removed so FullAWSAccess will still apply
upvoted 1 times

 **devakram** 11 months ago

no, I've just tested it in my account now, and B is the true answer. Although there were inherited SCPs coming from root and env which still showed in the SCP page for that OU, after detaching the allow all SCP, I was denied access on any other API except EC2.
upvoted 2 times

 **thanhnv142** 1 year, 1 month ago

B is correct: SCP have allow statement and this matches
upvoted 2 times

 **sarlos** 1 year, 1 month ago

a is the answer
upvoted 1 times

 **1123lluu** 1 year, 3 months ago

should be B, see example in here: <https://aws.amazon.com/blogs/security/get-more-out-of-service-control-policies-in-a-multi-account-environment/>
upvoted 1 times

 **zolthar_z** 1 year, 3 months ago

Selected Answer: A

Answer is A: You can't remove heritage policy from child OU
upvoted 2 times

 **learnwithaniket** 1 year, 3 months ago

Selected Answer: B

B is the right answer.


For a permission to be allowed for a specific account, there must be an explicit Allow statement at every level from the root through each OU in the direct path to the account (including the target account itself). This is why when you enable SCPs, AWS Organizations attaches an AWS managed SCP policy named FullAWSAccess which allows all services and actions. If this policy is removed and not replaced at any level of the organization, all OUs and accounts under that level would be blocked from taking any actions.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html
upvoted 3 times

 **tatarai1964** 1 year, 4 months ago

Selected Answer: B

"SCP evaluation follows a deny-by-default model, meaning that any permissions not explicitly allowed in the SCPs are denied. If an allow statement is not present, the default is deny." https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html#:~:text=SCP%20evaluation%20follows%20a%20deny-by-default%20model,meaning%20that%20any%20permissions%20not%20explicitly%20allowed%20in%20the%20SCPs%20are%20denied.%20If%20an%20allow%20statement%20is%20not%20present,%20the%20default%20is%20deny.
upvoted 4 times

 **jdxx000** 1 year, 5 months ago

Selected Answer: A

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_inheritance_mgmt.html
upvoted 1 times

 **zain1258** 1 year, 3 months ago

This URL does not explain the SCP.
upvoted 1 times


 **Radeeka** 1 year, 6 months ago

Selected Answer: A

Even the default policy is removed, Child OU will inherit the SCP from the Environment OU, which is AWSFullAccess. So the Child OU will still have full access.
upvoted 3 times

 **Gathix444** 1 year, 6 months ago

It's A, the new policy is an allow policy not deny, thus all permissions are granted to Dev OU.
upvoted 2 times

 **ixdb** 1 year, 6 months ago

Selected Answer: B

SCP can define an allow list – actions are prohibited by default, and you specify what services and actions are allowed.

upvoted 3 times

  **vhorman** 1 year, 7 months ago

Selected Answer: A

A is correct.

Development OU will inherit FullAccess from the Environments OU

no explicit DENY in the new AllowAllEc2 Policy

upvoted 4 times



  **Aja1** 1 year, 7 months ago

The answer is B.

When a policy is removed from an OU, the default policy for the parent OU is inherited. In this case, the default policy for the Environments OU is FullAWSAccess, which allows all API actions on all resources.



When the DevOps engineer replaces the FullAWSAccess policy with a policy that allows all actions on Amazon EC2 resources, the new policy will take precedence over the default policy. This means that all users in the Development OU will be allowed all API actions on EC2 resources. All other API actions will be denied.

upvoted 3 times

  **Gathix444** 1 year, 6 months ago

The last part is wrong. SCP doesnt deny anything unless you explicit define it.

upvoted 1 times

  **yorkicurke** 1 year, 3 months ago

because SCPs define the maximum permissions for an organization or organizational unit (OU) in AWS Organizations.

If an SCP doesn't explicitly grant permissions for an action, then that action is implicitly denied.


upvoted 1 times

  **yorkicurke** 1 year, 3 months ago

link;

https://repost.aws/questions/QUSHz1PpiJTOqWRuguGn_Trw/resource-and-iam-policy-with-scp

upvoted 1 times

  **FunkyFresco** 1 year, 8 months ago

Selected Answer: B

B is the correct option.

upvoted 4 times

A company is examining its disaster recovery capability and wants the ability to switch over its daily operations to a secondary AWS Region. The company uses AWS CodeCommit as a source control tool in the primary Region.

A DevOps engineer must provide the capability for the company to develop code in the secondary Region. If the company needs to use the secondary Region, developers can add an additional remote URL to their local Git configuration.

Which solution will meet these requirements?

- A. Create a CodeCommit repository in the secondary Region. Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repository. Create an AWS Lambda function that invokes the CodeBuild project. Create an Amazon EventBridge rule that reacts to merge events in the primary Region's CodeCommit repository. Configure the EventBridge rule to invoke the Lambda function.
- B. Create an Amazon S3 bucket in the secondary Region. Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucket. Create an AWS Lambda function that initiates the Fargate task. Create an Amazon EventBridge rule that reacts to merge events in the CodeCommit repository. Configure the EventBridge rule to invoke the Lambda function.
- C. Create an AWS CodeArtifact repository in the secondary Region. Create an AWS CodePipeline pipeline that uses the primary Region's CodeCommit repository for the source action. Create a cross-Region stage in the pipeline that packages the CodeCommit repository contents and stores the contents in the CodeArtifact repository when a pull request is merged into the CodeCommit repository.
- D. Create an AWS Cloud9 environment and a CodeCommit repository in the secondary Region. Configure the primary Region's CodeCommit repository as a remote repository in the AWS Cloud9 environment. Connect the secondary Region's CodeCommit repository to the AWS Cloud9 environment.

Suggested Answer: B

Community vote distribution



🗨️ 👤 **zijo** 8 months, 1 week ago

Why is not D a solution? If the developers in the secondary region can configure primary region's codecommit repository as a remote repository in the AWS Cloud9 environment they can do development and do all git functions remote.

upvoted 1 times

🗨️ 👤 **Gomer** 8 months, 3 weeks ago

A: (NO) "Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repository."

CodeBuild doesn't have the ability to do a "git mirror" operation itself. All online examples have CodeCommit actions calling Lambda (directly or through EventWatch) which calls fargate (or EC2) which does the actual git mirror

A: (NO) "Create an AWS Lambda function that invokes the CodeBuild project.

The is exactly the reverse from online examples

B: (NO) "Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucket."

Does it really make sense to use a "git" mirror operations copy from a CodeCommit repo to an S3 bucket? All online examples using "git" "mirror" have CodeCommit repo as remote target.

upvoted 1 times

🗨️ 👤 **Gomer** 8 months, 3 weeks ago

The specific requirement here isn't "disaster recovery capability" and "ability to switch over its daily operations to a secondary AWS Region." That is just being investigated.

The specific requirement is to "provide the capability for the company to develop code in the secondary Region."

"If the company needs to use the secondary Region, developers can add an additional remote URL to their local Git configuration."



To me this sounds to like the specific requirement here is only to provide developers with a complete remote development environment (not

to provide a DR solution)

If that is true, then using Cloud9 web development environment (includes git, etc.) with same local CodeCommit repo is acceptable


I'm not a developer, but the specific criteria wording and logic make me lean towards "C"

upvoted 1 times

  **Gomer** 8 months, 3 weeks ago



Actually, I meant to say I lean towards "D" (using Cloud9 as remote development environment)

upvoted 2 times

  **Gomer** 8 months, 3 weeks ago

Also want to add that "D" would work fine if you presume that the "git" "mirror" is also being done (though additional undefined step in the solution). Nothing says "D" is the complete solution. The ONLY requirement here is to provide developers a remote environment to develop in.

upvoted 1 times

  **Gomer** 8 months, 3 weeks ago

Flow:

AWS Example: CodeCommit(action) -----> Lambda -> Fargate task ("git clone --mirror" local repo, "git remote set-url --push origin" destination repo) -> CodeCommit(remote repo)

Solution "B": CodeCommit(action) -> EventBridge -> Lambda -> Fargate task ("git clone --mirror" local repo, "git remote set-url --push origin" destination repo) -> S3(remote bucket)

References:

<https://aws.amazon.com/blogs/devops/replicate-aws-codecommit-repository-between-regions-using-aws-fargate/>

<https://aws.amazon.com/cloud9/>

upvoted 2 times

  **thanhnv142** 1 year, 1 month ago

A is correct: <A DevOps engineer must provide the capability for the company to develop code in the secondary Region> means code commit

upvoted 3 times

  **thanhnv142** 1 year, 1 month ago

A is correct: < develop code in the secondary Region>: code commit cannot automatically clone cross-region. Must use a tool to do this duplication task

B: Using S3 as a secondary repo is incorrect

C and D: no mention of using codecommit as the secondary repo

upvoted 2 times

  **yuliaqwerty** 1 year, 2 months ago

Selected Answer: A

Agree answer is A

upvoted 2 times

  **svjl** 1 year, 3 months ago

B- It does the replication out of the box and meets the requirements


<https://aws.amazon.com/blogs/devops/replicate-aws-codecommit-repository-between-regions-using-aws-fargate/>

upvoted 2 times

  **bnagaraja9099** 1 year, 2 months ago

This part of B is incorrect. - It should use Code commit instead of S3. "Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucket. "



upvoted 1 times

  **vandergun** 1 year, 3 months ago

Selected Answer: A

A is at least operation and cost

upvoted 2 times

  **Dushank** 1 year, 5 months ago

Selected Answer: A



This solution meets all of the company's requirements:

It allows developers to add an additional remote URL to their local Git configuration to develop code in the secondary Region.

It is automated: the EventBridge rule will automatically invoke the Lambda function whenever a merge event occurs in the primary Region's CodeCommit repository.

It is reliable: the CodeBuild project will use Git to ensure that a perfect copy of the primary Region's CodeCommit repository is created in the secondary Region.


upvoted 4 times

  **RVivek** 1 year, 5 months ago

Selected Answer: A

<https://dev.to/apatil88/replicate-aws-codecommit-repositories-between-regions-using-codebuild-and-codepipeline-5fh1>

upvoted 4 times

  **ixdb** 1 year, 6 months ago

Selected Answer: B

B is right. <https://aws.amazon.com/cn/blogs/devops/replicate-aws-codecommit-repository-between-regions-using-aws-fargate/>

upvoted 2 times

  **zendevloper** 1 year, 5 months ago

B is wrong because it uses S3. Developers need a valid git remote URL.

Correct answer is A

upvoted 3 times

A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.

Which solution will meet these requirements?

- A. Use a buildspec file in AWS CodeBuild to restore the DB cluster from a snapshot of the production database, run integration tests, and drop the restored database after verification.
- B. Deploy the application to production. Configure an audit log of data control language (DCL) operations to capture database activities to perform if verification fails.
- C. Create a snapshot of the DB cluster before deploying the application. Use the Update requires:Replacement property on the DB instance in AWS CloudFormation to deploy the application and apply the changes.
- D. Ensure that the DB cluster is a Multi-AZ deployment. Deploy the application with the updates. Fail over to the standby instance if verification fails.

Suggested Answer: D

Community vote distribution

A (100%)

 **Ramdi1** Highly Voted 1 year, 1 month ago

Selected Answer: A

A is the solution which will allow testing without any such consequence
upvoted 5 times

 **jamesf** Most Recent 7 months, 2 weeks ago

Selected Answer: A

A correct as allow testing before real deployment.
<https://aws.amazon.com/blogs/devops/enhancing-automated-database-continuous-integration-with-aws-codebuild-and-amazon-rds-database-snapshot/>
upvoted 2 times

 **thanhnv142** 1 year, 1 month ago

Selected Answer: A

A is correct: This option allow testing before real deployment
B: < Deploy the application to production> : this would not allow testing before changes are made
C: <Create a snapshot of the DB cluster before deploying the application>: This means the same as B - would not allow testing before changes are made
D: <Ensure that the DB cluster is a Multi-AZ deployment. Deploy the application with the updates> - This deploy the app before testing, so it is incorrect
upvoted 4 times

 **thanhnv142** 1 year, 1 month ago

A is correct: <The DevOps team uses continuous integration to periodically verify that the application works> and <The DevOps team needs to test the changes before the changes are deployed to the production database> means codebuild
upvoted 1 times


 **Dushank** 1 year, 5 months ago

This solution meets all of the company's requirements:

It allows the DevOps team to test the changes before they are deployed to the production database.

It is automated: the CodeBuild buildspec file will automatically restore the DB cluster from a snapshot, run the integration tests, and drop the restored database after verification.

It is reliable: the CodeBuild buildspec file will ensure that the integration tests are run against a copy of the production database.
upvoted 1 times

 **ixdb** 1 year, 6 months ago

Selected Answer: A

A is right. All others will change the prod db.

upvoted 4 times

  **traveller37** 1 year, 6 months ago

I think it is A

<https://aws.amazon.com/blogs/devops/enhancing-automated-database-continuous-integration-with-aws-codebuild-and-amazon-rds-database-snapshot/>

upvoted 2 times

A company manages a multi-tenant environment in its VPC and has configured Amazon GuardDuty for the corresponding AWS account. The company sends all GuardDuty findings to AWS Security Hub.


Traffic from suspicious sources is generating a large number of findings. A DevOps engineer needs to implement a solution to automatically deny traffic across the entire VPC when GuardDuty discovers a new suspicious source.

Which solution will meet these requirements?

- A. Create a GuardDuty threat list. Configure GuardDuty to reference the list. Create an AWS Lambda function that will update the threat list. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.
- B. Configure an AWS WAF web ACL that includes a custom rule group. Create an AWS Lambda function that will create a block rule in the custom rule group. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.
- C. Configure a firewall in AWS Network Firewall. Create an AWS Lambda function that will create a Drop action rule in the firewall policy. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.
- D. Create an AWS Lambda function that will create a GuardDuty suppression rule. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.

Suggested Answer: B

Community vote distribution



traveller37 Highly Voted 1 year, 6 months ago

I think C:

<https://aws.amazon.com/blogs/security/automatically-block-suspicious-traffic-with-aws-network-firewall-and-amazon-guardduty/>
upvoted 14 times

traveller37 1 year, 6 months ago

Sorry i means B

upvoted 1 times

dencccc 1 year, 4 months ago

You mean C?

upvoted 1 times

RVivek Highly Voted 1 year, 5 months ago

Selected Answer: C

C is correct . Only Network Firewall can block traffic at VPC level.

A only updates the list , no blocking action

B- WAF and Web ACL can block only HTTPS traffic for a API/VPC endpoint/ Cloudfront distribution not for enire VPC

upvoted 11 times

jamesf Most Recent 7 months, 2 weeks ago

Selected Answer: C

C, AWS Network Firewall can block traffic at VPC level.

<https://aws.amazon.com/blogs/security/automatically-block-suspicious-traffic-with-aws-network-firewall-and-amazon-guardduty/>
upvoted 2 times

zijo 8 months, 1 week ago

Selected Answer: C

B blocks traffic at the http/https web traffic layer not for VPC layer

upvoted 1 times

thanhnv142 1 year, 1 month ago

Selected Answer: C

C is correct: <a solution to automatically deny traffic> means network FW.

A: irrelevant

B: We need network fw, not WAF

D: irrelevant

upvoted 3 times

🗨️ **yorkicurke** 1 year, 3 months ago

hmmm

is this the last question as of now(25th Nov 23)

upvoted 1 times

🗨️ **Dushank** 1 year, 5 months ago

Selected Answer: C

Here's the rationale for choosing this option:

AWS Network Firewall:

AWS Network Firewall is designed to provide centralized network traffic inspection and filtering. It's a suitable choice for implementing network-level controls.

Lambda Function for Automation:

Creating a Lambda function to trigger the creation of a Drop action rule in the firewall policy allows for automated response based on Security Hub findings. This enables you to take immediate action when suspicious sources are detected.

Specific Action (Drop):

The Drop action rule is effective for denying traffic from suspicious sources, effectively controlling access and preventing unwanted traffic.

This approach aligns well with the requirement to automatically deny traffic when GuardDuty identifies a new suspicious source, enhancing security in the multi-tenant VPC environment.

upvoted 6 times

🗨️ **RVivek** 1 year, 5 months ago

Selected Answer: B

A only will update threat list. the requirement is to block the traffic.

B is correct. Also it is event driven immediate action

upvoted 1 times

🗨️ **vladik820** 1 year, 5 months ago

Selected Answer: A

A is right

upvoted 1 times

A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.


A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.

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

- A. Update the default KMS key for Secrets Manager to allow only the Lambda function's execution role to decrypt
- B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt. Update Secrets Manager to use the new customer managed key
- C. Create a KMS customer managed key that trusts Secrets Manager and allows the account's root principal to decrypt. Update Secrets Manager to use the new customer managed key
- D. Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret
- E. Remove all KMS permissions from the Lambda function's execution role

Suggested Answer: CE

Community vote distribution



thanhnv142 Highly Voted 1 year, 1 month ago

Selected Answer: BD

B and D are correct: <update the infrastructure to ensure that only the Lambda function's execution role> means we need to ensure that lambda's IAM role has sufficient permissions and KMS policy allows Lambda's IAM role

A: cannot update default key

C: <allows the account's root principal to decrypt> this against the principal of least privilege

E: irrelevant

upvoted 5 times

heff_bezos Most Recent 5 months, 2 weeks ago

Selected Answer: BD

If default keys are the same as the AWS managed keys, then the answer is B. You cannot modify the "default" key's policy to allow access only from the Lambda execution role.

upvoted 1 times

jamesf 7 months, 2 weeks ago

Selected Answer: BD

I go for BD

upvoted 2 times

4555894 1 year ago

Selected Answer: BD

The requirement is to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.

upvoted 2 times

hotblooded 1 year, 1 month ago

Selected Answer: AD

```
{
  "Version": "2012-10-17",
  "Id": "key-consolepolicy-2",
  "Statement": [
```



```
{
  "Sid": "Allow use of the key",
  "Effect": "Allow",
  "Principal": {"AWS": [
    "arn:aws:iam::111122223333:role/KeyCreatorRole"
  ]},
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "here arn of secret manager"
}
```

I think A is correct answer , why to create CMK as customer is using default KMS
upvoted 1 times

  **zoltar_z** 1 year, 3 months ago

Selected Answer: BD

I think B:D
upvoted 4 times

  **radev** 1 year, 3 months ago

Selected Answer: BD

B, D

A is incorrect because updating the default KMS key for Secrets Manager to allow only the Lambda function's execution role to decrypt would grant access to all other resources using the default key, which violates the principle of least privilege.

C is incorrect because allowing the account's root principal to decrypt the secret would grant unnecessary access to the secret, which violates the principle of least privilege.



E is incorrect because removing all KMS permissions from the Lambda function's execution role would prevent the Lambda function from decrypting the secret, which is required for it to function properly.
upvoted 4 times

  **hotblooded** 1 year, 1 month ago

```
{
  "Version": "2012-10-17",
  "Id": "key-consolepolicy-2",
  "Statement": [
    {
      "Sid": "Allow use of the key",
      "Effect": "Allow",
      "Principal": {"AWS": [
        "arn:aws:iam::111122223333:role/KeyCreatorRole"
      ]},
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey"
      ],
      "Resource": "here arn of secret manager"
    }
  ]
}
```

```
}
```

I think A is correct answer , why to create CMK as customer is using default KMS
upvoted 1 times

  **hotblooded** 1 year, 1 month ago
Or we can ad below condition also

```
"Condition": {  
  "StringEquals": {  
    "kms:CallerAccount": "111122223333",  
    "kms:ViaService": "secretsmanager.us-west-2.amazonaws.com"  
  }  
}
```

upvoted 1 times

  **vandergun** 1 year, 3 months ago

Selected Answer: BD

I vote B,D

upvoted 2 times

A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance.

During testing, a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.


The DevOps engineer needs to prevent the loss of notification messages in the future.

Which solutions will meet this requirement? (Choose two.)

- A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.
- B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.
- C. Configure an Amazon Simple Queue Service (Amazon SQS) dead-letter queue for the SNS topic.
- D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic. Configure the Lambda function to process messages from the SQS queue.
- E. Replace the SNS topic with an Amazon EventBridge event bus. Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Suggested Answer: AD

Community vote distribution



vandergun Highly Voted 1 year, 3 months ago

Selected Answer: CD

The two solutions that will meet the requirement of preventing the loss of notification messages in the future are:

D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic. Configure the Lambda function to process messages from the SQS queue.

This solution will ensure that notification messages are delivered to the SQS queue even if the Lambda function is unavailable or the RDS DB instance is down. The Lambda function can then process the messages from the SQS queue at its own pace.

C. Configure an Amazon Simple Queue Service (Amazon SQS) dead-letter queue for the SNS topic.

This solution will ensure that notification messages that cannot be delivered to the RDS DB instance are not lost. Instead, they will be moved to a dead-letter queue. The DevOps engineer can then manually process the messages from the dead-letter queue.

upvoted 11 times

zolthar_z Highly Voted 1 year, 3 months ago

Selected Answer: CD

C:D, D is a best practice for this scenario, C because you can send failed SNS o SQS Dead letter queue, <https://docs.aws.amazon.com/sns/latest/dg/sns-dead-letter-queues.html>

upvoted 5 times

ce0df07 Most Recent 1 month ago

Selected Answer: BD

C does not make sense. Adding a DLQ to an SNS topic will capture messages that failed to be sent to the consumer, due to an issue with SNS. We are looking to capture messages that couldn't be processed due to an issue at the consumer side (RDS database to be exact, but the Lambda function will be aware of this).

Hence there are two solutions:

1. configure a SQS between the SNS topic and the Lambda function
2. configure a DLQ as a destination for failure for the Lambda function (raise an exception in the Lambda code)

upvoted 1 times

🗨️ 👤 **CHRIS12722222** 2 months, 3 weeks ago

Selected Answer: BD

Correct answer

<https://www.youtube.com/watch?v=rYFAdRCibyc>

upvoted 1 times

🗨️ 👤 **weixing** 5 months ago

BD

C. Dead-letter queues can only be added to SNS subscriptions, not to topics.

upvoted 1 times

🗨️ 👤 **h432ng** 7 months, 3 weeks ago

AD.

C is wrong, "Configuring an Amazon SNS dead-letter queue for a subscription" not for SNS topic

A is correct, with Dynamodb, admin can no longer "accidentally shut down the DB instance."

A fixes the root cause. With D an SQS is there, no need for DLQ for SNS. If lambda process data from SQS, what is SNS DLQ help here?

upvoted 2 times

🗨️ 👤 **thanhnv142** 1 year, 1 month ago

Selected Answer: CD

C and D are correct: <. While the database was down the company lost several of the SNS notification messages that were delivered during that time> means dead-letter queue in SQS and output SNS to SQS to store dead-letter queue

upvoted 4 times

🗨️ 👤 **zain1258** 1 year, 3 months ago

Selected Answer: BD

B & D are correct

upvoted 1 times

🗨️ 👤 **Gomer** 8 months, 2 weeks ago

Here's what I get when you break it down graphically between CD and BC:

CD: SNS > SQS(DLQ) > Lambda > RDS

BD: SNS > SQS > Lambda > SQS > RDS

The DLQ is just there to handle any SNS messages that have errors and can't be processed. There is no way you want/need two SQS queues in series (on either side of the Lambda). The ONLY thing you need to add for the requirements is queue to hold stuff while DB is down. The DLQ just makes sure even a messed up message data is captured for later review. Only C&D make any sense here.

upvoted 2 times

A company has an application that runs on Amazon EC2 instances. The company uses an AWS CodePipeline pipeline to deploy the application into multiple AWS Regions. The pipeline is configured with a stage for each Region. Each stage contains an AWS CloudFormation action for each Region.


When the pipeline deploys the application to a Region, the company wants to confirm that the application is in a healthy state before the pipeline moves on to the next Region. Amazon Route 53 record sets are configured for the application in each Region. A DevOps engineer creates a Route 53 health check that is based on an Amazon CloudWatch alarm for each Region where the application is deployed.

What should the DevOps engineer do next to meet the requirements?

- A. Create an AWS Step Functions workflow to check the state of the CloudWatch alarm. Configure the Step Functions workflow to exit with an error if the alarm is in the ALARM state. Create a new stage in the pipeline between each Region deployment stage. In each new stage, include an action to invoke the Step Functions workflow.
- B. Configure an AWS CodeDeploy application to deploy a CloudFormation template with automatic rollback. Configure the CloudWatch alarm as the instance health check for the CodeDeploy application. Remove the CloudFormation actions from the pipeline. Create a CodeDeploy action in the pipeline stage for each Region.
- C. Create a new pipeline stage for each Region where the application is deployed. Configure a CloudWatch alarm action for the new stage to check the state of the CloudWatch alarm and to exit with an error if the alarm is in the ALARM state
- D. Configure the CloudWatch agent on the EC2 instances to report the application status to the Route 53 health check. Create a new pipeline stage for each Region where the application is deployed. Configure a CloudWatch alarm action to exit with an error if the CloudWatch alarm is in the ALARM state.

Suggested Answer: D

Community vote distribution



🗳️ **heff_bezos** 5 months, 2 weeks ago

Selected Answer: A

There are no such things as cloudwatch alarm actions. The only things alarms can do is send notifications to an SNS topic. You can perform actions by using EventBridge (CloudWatch Log events) or Step Functions.

upvoted 2 times

🗳️ **zijo** 8 months ago

Selected Answer: D

D seems to be simple solution for me

The CloudWatch agent on EC2 instances can be configured to report the application status, and this information can then be used by Route 53 health checks.

Create Route 53 health checks that are based on the CloudWatch alarms. When you create a health check in Route 53, you can specify that the health check should be based on the state of a CloudWatch alarm. Route 53 health checks can be configured to treat the CloudWatch alarm state as Healthy or Unhealthy.

upvoted 1 times

🗳️ **govindr** 1 year, 1 month ago

D is correct - <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/monitoring-cloudwatch.html>

upvoted 1 times

🗳️ **thanhnv142** 1 year, 1 month ago

A is correct: <confirm that the application is in a healthy state before the pipeline moves on to the next Region.> means we need a new stage

B and C: no mention of creating a new stage

D: irrelevant

upvoted 3 times

🗳️ **a54b16f** 1 year, 1 month ago

Selected Answer: A

Exact scenario for Step usage: different routing options based on choices

upvoted 4 times

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

Selected Answer: A

A: <https://dev.to/aws-builders/dynamic-build-orchestration-using-codepipeline-codebuild-and-step-functions-2kpa>

upvoted 4 times

🗨️ 👤 **zain1258** 1 year, 3 months ago

Selected Answer: A

A is correct answer

upvoted 3 times

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

Selected Answer: A

A - 'If the state machine execution reaches a terminal status of FAILED, TIMED_OUT, or ABORTED, the action execution fails.'

<https://docs.aws.amazon.com/codepipeline/latest/userguide/action-reference-StepFunctions.html>

Can't be D because you can't update a Route53 healthcheck via the Cloudwatch agent

upvoted 3 times

A company plans to use Amazon CloudWatch to monitor its Amazon EC2 instances. The company needs to stop EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. The company must evaluate the metric every hour. The EC2 instances must continue to run if there is missing data for the NetworkPacketsIn metric during the evaluation period.

A DevOps engineer creates a CloudWatch alarm for the NetworkPacketsIn metric. The DevOps engineer configures a threshold value of 5 and an evaluation period of 1 hour.

Which set of additional actions should the DevOps engineer take to meet these requirements?

- A. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as breaching the threshold. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.
- B. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as not breaching the threshold. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
- C. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the threshold. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
- D. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as not breaching the threshold. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Suggested Answer: C

Community vote distribution

B (100%)

🗳️ **zolthar_z** Highly Voted 1 year, 3 months ago

Selected Answer: B

B: This is the reason <https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/UsingAlarmActions.html#AddingStopActions>
upvoted 7 times

🗳️ **zijo** Most Recent 6 months, 3 weeks ago

In CloudWatch alarms, datapoints are the individual metric values collected during each period. Here the period is 1 hour and 1 datapoint every hour collected. So 3 datapoints out of 12 is the alarm state because the network threshold has to be less than 5 for atleast 3 hours to Alarm state. The DevOps Engineer sets evaluation for every hour to look for the threshold value of 5. So if 1 hour has no data it is not breaching threshold. If there The alarm evaluates these datapoints against the conditions you set to determine whether it should trigger an action which is stopping an EC2 instance. My explanation for B
upvoted 1 times

🗳️ **PrasannaBalaji** 1 year, 2 months ago

Selected Answer: B

B is correct
upvoted 4 times

🗳️ **tom_cat** 1 year, 3 months ago

Selected Answer: B

I think B
upvoted 3 times

🗳️ **vandergun** 1 year, 3 months ago

Selected Answer: B

B should be corrected
upvoted 4 times

🗳️ **vandergun** 1 year, 3 months ago

B should be corrected
upvoted 3 times

A company manages 500 AWS accounts that are in an organization in AWS Organizations. The company discovers many unattached Amazon Elastic Block Store (Amazon EBS) volumes in all the accounts. The company wants to automatically tag the unattached EBS volumes for investigation.

A DevOps engineer needs to deploy an AWS Lambda function to all the AWS accounts. The Lambda function must run every 30 minutes to tag all the EBS volumes that have been unattached for a period of 7 days or more.

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

- A. Configure a delegated administrator account for the organization. Create an AWS CloudFormation template that contains the Lambda function. Use CloudFormation StackSets to deploy the CloudFormation template from the delegated administrator account to all the member accounts in the organization. Create an Amazon EventBridge event bus in the delegated administrator account to invoke the Lambda function in each member account every 30 minutes.
- B. Create a cross-account IAM role in the organization's member accounts. Attach the AWSLambda_FullAccess policy and the AWSCloudFormationFullAccess policy to the role. Create an AWS CloudFormation template that contains the Lambda function and an Amazon EventBridge scheduled rule to invoke the Lambda function every 30 minutes. Create a custom script in the organization's management account that assumes the role and deploys the CloudFormation template to the member accounts.
- C. Configure a delegated administrator account for the organization. Create an AWS CloudFormation template that contains the Lambda function and an Amazon EventBridge scheduled rule to invoke the Lambda function every 30 minutes. Use CloudFormation StackSets to deploy the CloudFormation template from the delegated administrator account to all the member accounts in the organization.
- D. Create a cross-account IAM role in the organization's member accounts. Attach the AmazonS3FullAccess policy and the AWSCodeDeployDeployerAccess policy to the role. Use AWS CodeDeploy to assume the role to deploy the Lambda function from the organization's management account. Configure an Amazon EventBridge scheduled rule in the member accounts to invoke the Lambda function every 30 minutes.

Suggested Answer: A

Community vote distribution

C (100%)

 **thanhnv142** Highly Voted 7 months ago

C is correct: <The Lambda function must run every 30 minutes to tag all the EBS volumes>: we should use a combination of eventbridge and Lambda

A: <. Create an Amazon EventBridge event bus in the delegated administrator account to invoke the Lambda function>: event bridge should be in each member account to monitor event, not in the delegated admin's account

B and D: These options create an IAM role in every member account, which is incorrect
upvoted 5 times

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

Selected Answer: C

C make the most sense
upvoted 3 times

 **a54b16f** 8 months ago

Selected Answer: C

NOT A: you don't want to run it for every user accounts
upvoted 3 times



 **yuliaqwerty** 8 months ago

Selected Answer: C

Agree with C
upvoted 3 times

 **davdan99** 8 months ago



Why no A?
upvoted 2 times

  **zain1258** 9 months, 2 weeks ago

Selected Answer: C

C is correct



upvoted 3 times

  **tom_cat** 9 months, 3 weeks ago

Selected Answer: C

C makes sense

upvoted 2 times

  **vandergun** 9 months, 3 weeks ago

Selected Answer: C

I vote C

upvoted 2 times

A company's production environment uses an AWS CodeDeploy blue/green deployment to deploy an application. The deployment includes Amazon EC2 Auto Scaling groups that launch instances that run Amazon Linux 2.

A working `appspec.yml` file exists in the code repository and contains the following text:

```
version: 0.0
os: linux
files:
  - source: /
    destination: /var/www/html/application
```

A DevOps engineer needs to ensure that a script downloads and installs a license file onto the instances before the replacement instances start to handle request traffic. The DevOps engineer adds a `hooks` section to the `appspec.yml` file.

Which hook should the DevOps engineer use to run the script that downloads and installs the license file?

- A. AfterBlockTraffic
- B. BeforeBlockTraffic
- C. BeforeInstall
- D. DownloadBundle

Suggested Answer: D

Community vote distribution

C (100%)

 **thanhnv142** Highly Voted 7 months ago

Selected Answer: C

C is correct: For blue/green deployment, Before install is one of several hooks that come before <the replacement instances> start to handle request traffic.

A and B: these hooks come after the replacement instances start to handle request traffic. They are hooks from the original instance, which are two of 3 last steps.

D: There is no such hook in blue/green deployment
upvoted 7 times

 **tom_cat** Highly Voted 9 months, 3 weeks ago

Selected Answer: C

A & B are not available for replacement instances - <https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#reference-appspec-file-structure-hooks-availability>

D - "Reserved for CodeDeploy operations. Cannot be used to run scripts."
upvoted 5 times

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

Selected Answer: C


C is correct
upvoted 2 times

 **twogyt** 7 months, 3 weeks ago

Selected Answer: C

is C: A DevOps engineer needs to ensure that a script downloads and "installs" a license file onto the instances "before" the replacement instances start to handle request traffic

upvoted 2 times

 **vandergun** 9 months, 3 weeks ago

Selected Answer: C

C should be correct
upvoted 3 times

A company has an application that includes AWS Lambda functions. The Lambda functions run Python code that is stored in an AWS CodeCommit repository. The company has recently experienced failures in the production environment because of an error in the Python code. An engineer has written unit tests for the Lambda functions to help avoid releasing any future defects into the production environment.

The company's DevOps team needs to implement a solution to integrate the unit tests into an existing AWS CodePipeline pipeline. The solution must produce reports about the unit tests for the company to view.

Which solution will meet these requirements?

- A. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create a buildspec.yml file in the CodeCommit repository. In the buildspec.yml file, define the actions to run a CodeGuru review.
- B. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create a CodeBuild report group. Create a buildspec.yml file in the CodeCommit repository. In the buildspec.yml file, define the actions to run the unit tests with an output of JUNITXML in the build phase section. Configure the test reports to be uploaded to the new CodeBuild report group.
- C. Create a new AWS CodeArtifact repository. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create an appspec.yml file in the original CodeCommit repository. In the appspec.yml file, define the actions to run the unit tests with an output of CUCUMBERJSON in the build phase section. Configure the tests reports to be sent to the new CodeArtifact repository.
- D. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create a new Amazon S3 bucket. Create a buildspec.yml file in the CodeCommit repository. In the buildspec.yml file, define the actions to run the unit tests with an output of HTML in the phases section. In the reports section, upload the test reports to the S3 bucket.

Suggested Answer: C

Community vote distribution

B (100%)

 **thanhnv142** Highly Voted 7 months ago

Selected Answer: B

B is correct: for unit test, we need codebuild

A: codeguru is for code analysis, not unit test

C: This option mentions pushing reports to CodeArtifact repository, which is incorrect


D: This option push reports to S3, which is incorrect. We should upload report to codebuild report group
upvoted 6 times

 **zain1258** Most Recent 9 months, 2 weeks ago

Selected Answer: B

B is correct

upvoted 3 times

 **KobraKai** 9 months, 2 weeks ago

I think B as per link:

<https://docs.aws.amazon.com/codebuild/latest/userguide/test-reporting.html>


upvoted 3 times

 **tom_cat** 9 months, 3 weeks ago

Selected Answer: B

I think it should be B

upvoted 3 times

 **vandergun** 9 months, 3 weeks ago

Selected Answer: B

B is corrected

upvoted 4 times

A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.

A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.

Which SCP will meet these requirements?

- A.
- ```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": "*",
 "Resource": "*",
 "Condition": {
 "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
 }
 }
]
}
```
- B.
- ```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Principal": { "AWS": "arn:aws:iam::*:root" }
    }
  ]
}
```
- C.
- ```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Deny",
 "Action": "*",
 "Resource": "*",
 "Condition": {
 "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
 }
 }
]
}
```
- D.
- ```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Principal": "root"
    }
  ]
}
```

Suggested Answer: B

Community vote distribution

C (100%)

 **tom_cat** Highly Voted 1 year, 3 months ago

Selected Answer: C

I believe it should be C

https://docs.aws.amazon.com/organizations/latest/userguide/best-practices_member-acct.html#bp_member-acct_use-scp

upvoted 8 times

 **thanhnv142** Highly Voted 1 year, 1 month ago

Selected Answer: C

C is correct: < will prevent the root user in member accounts> this means deny action

A and D: irrelevant (mention allow statement)

B: scp does not have principal element. only condition

upvoted 7 times

🗨️ 👤 **Gomer** Most Recent 8 months, 2 weeks ago

Selected Answer: C

A slightly more concise version of "C" is a "strongly recommended" control to deny root access in member accounts. See the example:

<https://docs.aws.amazon.com/controltower/latest/controlreference/strongly-recommended-controls.html#disallow-root-user-actions>

upvoted 1 times

🗨️ 👤 **c3518fc** 10 months, 3 weeks ago

Selected Answer: C

https://docs.aws.amazon.com/organizations/latest/userguide/best-practices_member-acct.html#bp_member-acct_use-scp

upvoted 1 times

🗨️ 👤 **DanShone** 12 months ago

Selected Answer: C

C is correct

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year ago

Selected Answer: C

C no debate

upvoted 2 times

🗨️ 👤 **manman7** 1 year, 3 months ago

It's C, based on the documentation :

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.html#example-scp-root-user

upvoted 4 times

🗨️ 👤 **zain1258** 1 year, 3 months ago

Selected Answer: C

C looks correct

upvoted 2 times

A company uses AWS and has a VPC that contains critical compute infrastructure with predictable traffic patterns. The company has configured VPC flow logs that are published to a log group in Amazon CloudWatch Logs.


The company's DevOps team needs to configure a monitoring solution for the VPC flow logs to identify anomalies in network traffic to the VPC over time. If the monitoring solution detects an anomaly, the company needs the ability to initiate a response to the anomaly.

How should the DevOps team configure the monitoring solution to meet these requirements?

- A. Create an Amazon Kinesis data stream. Subscribe the log group to the data stream. Configure Amazon Kinesis Data Analytics to detect log anomalies in the data stream. Create an AWS Lambda function to use as the output of the data stream. Configure the Lambda function to write to the default Amazon EventBridge event bus in the event of an anomaly finding.
- B. Create an Amazon Kinesis Data Firehose delivery stream that delivers events to an Amazon S3 bucket. Subscribe the log group to the delivery stream. Configure Amazon Lookout for Metrics to monitor the data in the S3 bucket for anomalies. Create an AWS Lambda function to run in response to Lookout for Metrics anomaly findings. Configure the Lambda function to publish to the default Amazon EventBridge event bus.
- C. Create an AWS Lambda function to detect anomalies. Configure the Lambda function to publish an event to the default Amazon EventBridge event bus if the Lambda function detects an anomaly. Subscribe the Lambda function to the log group.
- D. Create an Amazon Kinesis data stream. Subscribe the log group to the data stream. Create an AWS Lambda function to detect log anomalies. Configure the Lambda function to write to the default Amazon EventBridge event bus if the Lambda function detects an anomaly. Set the Lambda function as the processor for the data stream.

Suggested Answer: A

Community vote distribution



giovanna_mag Highly Voted 1 year, 2 months ago

Selected Answer: B

I think it's B, Amazon Lookout for metrics can detect anomalies from S3 bucket and trigger Lambda
<https://aws.amazon.com/lookout-for-metrics/>
 upvoted 7 times

ce0df07 Most Recent 1 month ago

Selected Answer: A

Using KDA for anomaly detection, means you can use the built-in RCF (Random Cut Forest) ML algorithm.
 Option B: Firehose and S3 adds latency. Further, Amazon Lookout for Metrics is more suitable for business metrics than network traffic.
 Options C&D require you to implement your own Lambda function, which means more error prone and more maintenance.
 upvoted 1 times

[Removed] 6 months, 3 weeks ago

Selected Answer: B

B for me
 upvoted 1 times

jamesf 7 months, 1 week ago

Selected Answer: B

- Data Streaming: Use Amazon Kinesis Data Firehose to deliver VPC flow logs from CloudWatch Logs to an Amazon S3 bucket.
- Anomaly Detection: Amazon Lookout for Metrics will monitor the data in the S3 bucket and automatically detect anomalies in the network traffic.
- Event Response: When Lookout for Metrics detects an anomaly, it triggers an AWS Lambda function. The Lambda function will then publish an event to the Amazon EventBridge event bus, which can further initiate automated responses, notifications, or alerts.

upvoted 4 times

trungtd 8 months ago

Selected Answer: B

Although option A uses Kinesis Data Analytics for anomaly detection, setting up and maintaining custom analytics and anomaly detection logic is more complex and less efficient compared to using a managed service like Lookout for Metrics.

upvoted 2 times

🗳️ 👤 **xdkonorek2** 8 months, 1 week ago

Selected Answer: B

A is wrong because kinesis data analytics output must be either kinesis data stream or firehose, can't be lambda directly so there is a missing component

upvoted 2 times

🗳️ 👤 **Gomer** 8 months, 2 weeks ago

I've reviewed most of the comments, and it seems like everyone is just repeating themselves. I've "googled" and looked at the references. I found examples of both kinesis data streams, kinesis data analytics and firehose. The one step in "A" I have a problem with is "Create an AWS Lambda function to use as the output of the data stream." How can Lambda be an output of a data stream "over time"? I don't think you can identify an anomaly "over time" unless you've got persistent storage for the data (which can be reparsed as necessary to compare past with present). I'm leaning towards "B" unless someone can convince me otherwise (and not by just repeating what others have already said).

upvoted 3 times

🗳️ 👤 **tsangckl** 8 months, 3 weeks ago

Selected Answer: A

Option B involves using Amazon Lookout for Metrics, which is not designed for real-time anomaly detection.

upvoted 2 times

🗳️ 👤 **Gomer** 8 months, 2 weeks ago

I see the "over time" requirement as implying some ability to parse the past with the present in order for ML to assess an anomaly. I don't see the words "real time" in the requirements. The "over time" requirement is not specific enough, but until there are more specifics, it would be reasonable to presume it means your trying to discover current anomalies by comparing traffic from against days, weeks or months ago.

upvoted 1 times

🗳️ 👤 **seetpt** 10 months, 1 week ago

Selected Answer: B

i think B

upvoted 2 times

🗳️ 👤 **c3518fc** 10 months, 3 weeks ago

Selected Answer: B

Lookout for Metrics automatically detects and diagnoses anomalies (outliers from the norm) in business and operational data. It's a fully managed ML service, which uses specialized ML models to detect anomalies based on the characteristics of your data. You don't need ML experience to use Lookout for Metrics.

Kinesis Data Analytics Studio provides an interactive notebook experience powered by Apache Zeppelin and Apache Flink to analyze streaming data. It also helps productionize your analytics application by building and deploying code as a Kinesis data analytics application straight from the notebook. <https://aws.amazon.com/blogs/machine-learning/smart-city-traffic-anomaly-detection-using-amazon-lookout-for-metrics-and-amazon-kinesis-data-analytics-studio/>

upvoted 3 times

🗳️ 👤 **stoy123** 11 months, 2 weeks ago

Selected Answer: A

A. If you google "detecting anomalies in vpc flow logs" every article suggests Kinesis Data Analytics

upvoted 1 times

🗳️ 👤 **CloudHandsOn** 11 months, 4 weeks ago

Selected Answer: A

I'll go with A. Mainly because Kinesis data analytics has anomaly detection using a random cut forest function:

<https://docs.aws.amazon.com/kinesisanalytics/latest/dev/app-anomaly-detection.html>

upvoted 2 times

🗳️ 👤 **DanShone** 12 months ago

Selected Answer: B

B - Amazon Lookout for Metrics Automatically detect anomalies within metrics and identify their root causes. So would fit the requirements

upvoted 3 times

🗳️ 👤 **ogerber** 12 months ago

Selected Answer: A

Option A is preferable for scenarios requiring real-time processing and anomaly detection in streaming data, such as VPC flow logs, with the capability to quickly initiate responses to detected anomalies. It offers a more streamlined and immediate approach to monitoring and responding to network traffic anomalies, making it highly suitable for the company's needs regarding their critical compute infrastructure with predictable traffic patterns.

Option B might still be considered if the company's workflow is more adapted to batch processing and the delays inherent in data delivery and processing are acceptable. However, for immediate anomaly detection and response, Option A stands out as the more appropriate solution.

upvoted 1 times

  **dzn** 1 year ago

Selected Answer: A

Kinesis Data Firehose determines how often to write to S3 by buffer settings, which is not realtime enough to handle VPC flow log, which can be fatal depending on the content of the `CRITICAL compute infrastructure`. Kinesis Data Analytics has machine learning solutions such as RANDOM_CUT_FOREST in addition to fixed detection by normal SQL.


upvoted 3 times

  **[Removed]** 1 year ago

Selected Answer: B

B without a doubt

upvoted 3 times

  **fdxxx** 1 year ago

Option B is the most suitable for the scenario.

Kinesis Data Firehose: It allows the streaming of data to an S3 bucket, providing a durable storage solution.

Lookout for Metrics: It is designed to detect anomalies in your data and can be configured to monitor the data stored in the S3 bucket for anomalies.

upvoted 3 times

AnyCompany is using AWS Organizations to create and manage multiple AWS accounts. AnyCompany recently acquired a smaller company, Example Corp. During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation. AnyCompany moved the new member account under an OU that is dedicated to Example Corp.

AnyCompany's DevOps engineer has an IAM user that assumes a role that is named OrganizationAccountAccessRole to access member accounts. This role is configured with a full access policy. When the DevOps engineer tries to use the AWS Management Console to assume the role in Example Corp's new member account, the DevOps engineer receives the following error message: "Invalid information in one or more fields. Check your information or contact your administrator."

Which solution will give the DevOps engineer access to the new member account?

- A. In the management account, grant the DevOps engineer's IAM user permission to assume the OrganizationAccountAccessRole IAM role in the new member account.
- B. In the management account, create a new SCP. In the SCP, grant the DevOps engineer's IAM user full access to all resources in the new member account. Attach the SCP to the OU that contains the new member account.
- C. In the new member account, create a new IAM role that is named OrganizationAccountAccessRole. Attach the AdministratorAccess AWS managed policy to the role. In the role's trust policy, grant the management account permission to assume the role.
- D. In the new member account, edit the trust policy for the OrganizationAccountAccessRole IAM role. Grant the management account permission to assume the role.

Suggested Answer: D

Community vote distribution

C (85%)


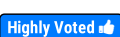
Other

 **radev**  1 year, 3 months ago

Selected Answer: C

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html#orgs_manage_accounts_create-cross-account-role

upvoted 6 times

 **thanhnv142**  1 year, 1 month ago

Selected Answer: C

C is correct: <assume the role in Example Corp's new member account> means this role has not been properly configured (or even not created)

A: only mention assuming the role, not create it.

B: scp has nothing to do here

D: only mention create trust relationship

upvoted 5 times

 **AC2021**  1 month, 2 weeks ago

Selected Answer: D


The role is already there. Why create a new one?

upvoted 1 times

 **ce0df07** 1 month ago

The role would not be automatically created in accounts that are added through invitation (as opposed to accounts created within the organization).

upvoted 1 times

 **Simba84** 2 months, 2 weeks ago

Selected Answer: D

Correct Answer is D

Role Trust Policy Issue:

When a new account is invited and joins an AWS Organization, the OrganizationAccountAccessRole is typically created automatically.

This role allows the management account to access member accounts, but its trust policy must explicitly grant the management account

permission to assume the role.

If this trust policy is not configured correctly, the management account cannot assume the role, leading to the error message.

upvoted 3 times



  **ce0df07** 1 month ago

Incorrect. See

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html#orgs_manage_accounts_create-cross-account-role

"When you create a member account using AWS Organizations, Organizations automatically creates an IAM role in the account that grants administrator access to the management account. For invited member accounts, you must manually create the role."

upvoted 1 times

  **eugene2owl** 3 months, 1 week ago

Selected Answer: C

I've spent like 30 mins, and now I've got the most full explanation.

Correct answer is "C"

While "D" is NOT FULLY describing what needs to be done (so it's wrong).

The thing you need to know to answer this question is the following:

* if account is generated (meaning NEW account CREATED) within the Org, then this account will automatically have a proper role

"OrganizationAccountAccessRole"

* if account is invited (meaning EXISTING account ADDED) to Org, then this account will NOT have such role

Question says, that Management Account tries to assume a role called "OrganizationAccountAccessRole" from member account, but it gets an error saying like "there is no such thing which you request".

So to fix an error you need:

1) Create a IAM Role "OrganizationAccountAccessRole" in a member account

2) Give it FullAccess Policy

3) Allow Management Account to assume this role via its Trust Relationship

upvoted 4 times

  **hamzaBennis** 3 months, 3 weeks ago

member accounts that you invite to join your organization do not automatically get an administrator role created.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_create-cross-account-role.html


upvoted 2 times

  **VerRi** 4 months ago

Selected Answer: D

"IAM user that assumes a role that is named OrganizationAccountAccessRole", the role is already there

upvoted 3 times

  **heff_bezos** 5 months, 2 weeks ago

Selected Answer: D

The question states that the role already exists with full access policy. This role exists in the new member account. We need the IAM user from the management account the ability to assume it.

upvoted 3 times

  **aws_god** 6 months ago

Selected Answer: D

This role is created by default in member accounts. See:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html

upvoted 2 times

  **heff_bezos** 5 months, 2 weeks ago

"By default, if you create a member account as part of your organization, AWS automatically creates a role in the account that grants administrator permissions to IAM users in the management account who can assume the role. By default, that role is named OrganizationAccountAccessRole"

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_create-cross-account-role.html

upvoted 1 times

🗳️ 👤 **c3518fc** 10 months, 3 weeks ago

Selected Answer: C

To create an AWS Organizations administrator role in a member account

Sign in to the IAM console at <https://console.aws.amazon.com/iam/>. You must sign in as an IAM user, assume an IAM role, or sign in as the root user (not recommended) in the member account. The user or role must have permission to create IAM roles and policies.

In the IAM console, navigate to Roles and then choose Create role.

Choose AWS account, and then select Another AWS account.

Enter the 12-digit account ID number of the management account that you want to grant administrator access to. Under Options, please note the following:

On the Add permissions page, choose the AWS managed policy named AdministratorAccess and then choose.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html#orgs_manage_accounts_create-cross-account-role

upvoted 4 times

🗳️ 👤 **Andy11234912** 11 months ago

Selected Answer: D

not c, the role is already created

upvoted 2 times

🗳️ 👤 **Jay_2pt0_1** 8 months, 3 weeks ago

From reading the question, I'm not sure.

upvoted 1 times

🗳️ 👤 **sirronido** 11 months, 1 week ago

D.. the role is already created, what is needed is just update the trust policy

upvoted 1 times

🗳️ 👤 **DanShone** 12 months ago

C is correct

upvoted 1 times

🗳️ 👤 **twogyt** 1 year, 1 month ago

Selected Answer: C

C is correct

upvoted 4 times

🗳️ 👤 **zain1258** 1 year, 3 months ago

Selected Answer: C

C is correct

upvoted 4 times

🗳️ 👤 **tom_cat** 1 year, 3 months ago

For invited accounts the OrganizationAccountAccessRole needs to be created:

Member accounts that you invite to join your organization do not automatically get an administrator role created. You have to do this manually, as shown in the following procedure. This essentially duplicates the role automatically set up for created accounts. We recommend that you use the same name, OrganizationAccountAccessRole, for your manually created roles for consistency and ease of remembering.

upvoted 2 times

🗳️ 👤 **tom_cat** 1 year, 3 months ago

So I believe it's C.

upvoted 2 times

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

Selected Answer: A

A should be correct

upvoted 1 times

A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.

Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.

During testing, the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.

Which solution will meet these requirements?

- A. Increase the retry attempts.
- B. Configure the setting to split the batch when an error occurs.
- C. Increase the concurrent batches per shard.
- D. Decrease the maximum age of record.

Suggested Answer: B

Community vote distribution

B (100%)

  **c3518fc**  10 months, 3 weeks ago

Selected Answer: B

When consuming records from a Kinesis data stream using AWS Lambda, the function can process records in batches. By default, if any record in the batch fails to process, the entire batch is sent to the dead-letter queue.

To avoid sending errorless records to the dead-letter queue, the Lambda function's event source options should be configured to split the batch when an error occurs. This setting is called `batchWindow` and can be configured in the event source mapping for the Lambda function.

When `batchWindow` is set to `TRIM_HORIZON`, the Lambda function will split the batch at the first record that causes an error and send only the failed records to the dead-letter queue. The remaining errorless records in the batch will continue to be processed by the function.

upvoted 11 times


  **Gomer** 8 months, 2 weeks ago

Seemingly very good explanation, though I had trouble finding any references other than this:

"BisectBatchOnFunctionError" "If the function returns an error, split the batch in two and retry. The default value is false."

`aws lambda update-event-source-mapping --bisect-batch-on-function-error [...]`

upvoted 1 times

  **zolthar_z**  1 year, 3 months ago

Selected Answer: B

B: <https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html#services-kinesis-eventsourcemapping>

upvoted 7 times

  **tinyshare**  7 months ago

B:


<https://aws.amazon.com/blogs/big-data/best-practices-for-consuming-amazon-kinesis-data-streams-using-aws-lambda/>

upvoted 1 times

  **thanhnv142** 1 year, 1 month ago

B is correct: <(Amazon SQS) dead-letter queue as an on-failure destination>: split the batch into 2 parts: success ones and error ones. error ones come to dead queue

upvoted 3 times

  **zain1258** 1 year, 3 months ago

Selected Answer: B

B is correct

upvoted 4 times

  **vandergun** 1 year, 3 months ago

Selected Answer: B

B is corrected

upvoted 5 times


A company has microservices running in AWS Lambda that read data from Amazon DynamoDB. The Lambda code is manually deployed by developers after successful testing. The company now needs the tests and deployments be automated and run in the cloud. Additionally, traffic to the new versions of each microservice should be incrementally shifted over time after deployment.

What solution meets all the requirements, ensuring the MOST developer velocity?

- A. Create an AWS CodePipeline configuration and set up a post-commit hook to trigger the pipeline after tests have passed. Use AWS CodeDeploy and create a Canary deployment configuration that specifies the percentage of traffic and interval.
- B. Create an AWS CodeBuild configuration that triggers when the test code is pushed. Use AWS CloudFormation to trigger an AWS CodePipeline configuration that deploys the new Lambda versions and specifies the traffic shift percentage and interval.
- C. Create an AWS CodePipeline configuration and set up the source code step to trigger when code is pushed. Set up the build step to use AWS CodeBuild to run the tests. Set up an AWS CodeDeploy configuration to deploy, then select the CodeDeployDefault.LambdaLinear10PercentEvery3Minutes option.
- D. Use the AWS CLI to set up a post-commit hook that uploads the code to an Amazon S3 bucket after tests have passed. Set up an S3 event trigger that runs a Lambda function that deploys the new version. Use an interval in the Lambda function to deploy the code over time at the required percentage.

Suggested Answer: A

Community vote distribution



PrasannaBalaji Highly Voted 8 months, 2 weeks ago

Selected Answer: C

Agree C is correct
upvoted 6 times

c3518fc Highly Voted 4 months, 2 weeks ago

Selected Answer: C

This solution provides the following benefits:

Automation: The entire process, from code push to testing and deployment, is automated, reducing manual effort and increasing developer velocity.

Integration: By using AWS CodePipeline, CodeBuild, and CodeDeploy, you leverage fully managed services that are designed to work together seamlessly.

Incremental Deployment: The CodeDeployDefault.LambdaLinear10PercentEvery3Minutes option ensures a smooth and controlled migration of traffic to the new versions of your microservices, minimizing the risk of downtime or disruption.

upvoted 5 times

DanShone Most Recent 5 months, 3 weeks ago

Selected Answer: C

C is correct
upvoted 3 times

Shasha1 6 months ago

C

There is no 'pre-commit' hook option in the Lambda deployment hook (Canary); only 'before allowing traffic' and 'after allowing traffic' options are available. Therefore, the 'LambdaLinear10PercentEvery3Minutes' option, which is a canary deployment method, enables a linear deployment strategy, gradually shifting traffic to the new versions at a rate of 10% every 3 minutes.

https://medium.com/@Da_vidgf/canary-deployments-in-serverless-applications-b0f47fa9b409

upvoted 3 times

Chelseajcole 7 months ago

if it is canary, why not a?
upvoted 1 times

thanhnv142 7 months ago

A is correct: canary deployment

upvoted 1 times

  **twogyt** 7 months, 3 weeks ago

Selected Answer: C

c is correct

upvoted 4 times

  **a54b16f** 8 months ago

Selected Answer: C

B is wrong, why would you trigger a pipeline when TEST code is pushed

upvoted 4 times

  **csG13** 8 months, 2 weeks ago

Selected Answer: C

Answer is C, canary deployment

upvoted 4 times

  **PrasannaBalaji** 8 months, 2 weeks ago

Selected Answer: B

B is correct

upvoted 1 times

A company is building a web and mobile application that uses a serverless architecture powered by AWS Lambda and Amazon API Gateway. The company wants to fully automate the backend Lambda deployment based on code that is pushed to the appropriate environment branch in an AWS CodeCommit repository.

The deployment must have the following:

- Separate environment pipelines for testing and production
- Automatic deployment that occurs for test environments only

Which steps should be taken to meet these requirements?

- A. Configure a new AWS CodePipeline service. Create a CodeCommit repository for each environment. Set up CodePipeline to retrieve the source code from the appropriate repository. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
- B. Create two AWS CodePipeline configurations for test and production environments. Configure the production pipeline to have a manual approval step. Create a CodeCommit repository for each environment. Set up each CodePipeline to retrieve the source code from the appropriate repository. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
- C. Create two AWS CodePipeline configurations for test and production environments. Configure the production pipeline to have a manual approval step. Create one CodeCommit repository with a branch for each environment. Set up each CodePipeline to retrieve the source code from the appropriate branch in the repository. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
- D. Create an AWS CodeBuild configuration for test and production environments. Configure the production pipeline to have a manual approval step. Create one CodeCommit repository with a branch for each environment. Push the Lambda function code to an Amazon S3 bucket. Set up the deployment step to deploy the Lambda functions from the S3 bucket.

Suggested Answer: C

Community vote distribution

C (100%)

 **thanhnv142**  7 months ago

Selected Answer: C

C is correct: <Separate environment pipelines for testing and production> means codepipeline <code that is pushed to the appropriate environment branch in an AWS CodeCommit repository.> means code CodeCommit

A: no mention of creating Separate env for test and dev

B: <Create a CodeCommit repository for each environment> should not do this. We should create a branch for each env

D: no mention of code pipelines

upvoted 6 times

 **c3518fc**  4 months, 2 weeks ago

Selected Answer: C

By creating two CodePipeline configurations, using a single CodeCommit repository with branches for each environment, and deploying Lambda functions with CloudFormation, this solution meets the requirements while following best practices for source code management, continuous delivery, and infrastructure as code.

upvoted 3 times

 **PrasannaBalaji** 8 months, 2 weeks ago

Selected Answer: C

C is correct

First, A&B both are in-correct: As a basic policy - do not create a repo for the same code for multiple environments. Always create a branch from the same repo. The strategy is wrong for A&B.

Now C&D: D uses Lambda function with s3, whereas C uses code pipeline to store and build. Using code pipeline is a smart choice rather than using S3 as a code pipeline that offers better branching strategy and controls. I will go with 'C'.

upvoted 3 times

 **csG13** 8 months, 2 weeks ago

Selected Answer: C

It's C - unique env and also distinct resources in aws codepipeline would result to pull from both repos on every update of either repo.

upvoted 2 times