Question #1 *Topic 1*

A company needs to architect a hybrid DNS solution. This solution will use an Amazon Route 53 private hosted zone for the domain cloud.example.com for the resources stored within VPCs.
The company has the following DNS resolution requirements:
On-premises systems should be able to resolve and connect to cloud.example.com.
All VPCs should be able to resolve cloud.example.com.
There is already an AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway.
Which architecture should the company use to meet these requirements with the HIGHEST performance?

A. Associate the private hosted zone to all the VPCs. Create a Route 53 inbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

B. Associate the private hosted zone to all the VPCs. Deploy an Amazon EC2 conditional forwarder in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the conditional forwarder.

C. Associate the private hosted zone to the shared services VPCreate a Route 53 outbound resolver in the shared services VPAttach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the outbound resolver.

D. Associate the private hosted zone to the shared services VPC. Create a Route 53 inbound resolver in the shared services VPC. Attach the shared services VPC to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

Suggested Answer: *D*

Community vote distribution

A (89%) 11%

robertohyena Highly Voted 9 months, 1 week ago

A. Correct answer. Source: https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-with-amazon-route-53-and-aws-transit-gateway/

NOT B. EC2 conditional forwarder will not meet Highest performance requirement.

NOT C. Missing: Need to associate private hosted zone to all VPC.
"All VPC's will need to associate their private hosted zones to all other VPC's if required to."
Source: https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-with-amazon-route-53-and-aws-transit-gateway/

NOT D. Missing: Need to associate private hosted zone to all VPC.
"All VPC's will need to associate their private hosted zones to all other VPC's if required to."
Source: https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-with-amazon-route-53-and-aws-transit-gateway/
upvoted 57 times

nokkosurdu 1 week, 4 days ago
here is valid source - https://appurl.io/xUl7PRUEfr
upvoted 1 times

awsylum 1 year, 4 months ago
In your link, you missed this sentence:
"The most reliable, performant and low-cost approach is to share and associate private hosted zones directly to all VPCs that need them." You share the PHZ via the Shared Services VPC. You use the .2 DNS Resolver Address in each VPC to connect to the PHZ in the shared services VPC for domain resolution.
upvoted 2 times

alexkro 1 year, 3 months ago
You forgot an additional condition mentioned in the question: "All VPCs should be able to resolve cloud.example.com." Nobody said there are only shared VPCs there.

upvoted 1 times

zhangyu20000 Highly Voted 2 years, 6 months ago

A because it requires all VPC can resolve the example.com. All VPCs must be associated with private hosted zone

upvoted 10 times

Robertwilliamm Most Recent 4 days, 3 hours ago

Selected Answer: A

A is Correct Option

Thanks to SkillCertExams I successfully cleared my SAP-C02 exam today.

upvoted 1 times

SofieneGh 2 months, 1 week ago

Selected Answer: A

A is correct

upvoted 1 times

CBMAN 2 months, 2 weeks ago

Selected Answer: A

A is correct

upvoted 1 times

diazed 2 months, 2 weeks ago

Selected Answer: A

A is correct.

upvoted 1 times

pekomari 3 months, 1 week ago

Selected Answer: A

529[illegible]80%[illegible]

upvoted 1 times

mssc 3 months, 3 weeks ago

Selected Answer: A

I have just taken the exam and 75 percent of the questions were from here. Prepare well for these Questions. Good Luck!

upvoted 3 times

FlyingHawk 4 months, 3 weeks ago

Selected Answer: A

On-premises systems should be able to resolve and connect to cloud.example.com, it is inbound resolver, C is incorrect.

All VPCS will need to associate their private zones to the Transit Gateway, associated only the shared VPC with TGW forces all the DNS query from other VPCS forward to shared VPC, add the latency. d is incorrect

upvoted 1 times

pk0619 6 months, 2 weeks ago

Selected Answer: A

When a Route 53 private hosted zone needs to be resolved in multiple VPCs and AWS accounts as described earlier, the most reliable pattern is to share the private hosted zone between accounts and associate it to each VPC that needs it.

upvoted 1 times

jrheen 7 months, 3 weeks ago

A. Correct answer. Source: https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-with-amazon-route-53-and-aws-transit-gateway/

upvoted 1 times

TariqKipkemei 8 months ago

Selected Answer: A

Associate the private hosted zone to all the VPCs. Create a Route 53 inbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

upvoted 1 times

to_to 8 months, 2 weeks ago

Selected Answer: D

1-1. Private hosted zone One Account -> 2 Account PHZ is not Equals.
1-2. VPCs in Private hosted zone
2. On-Premise -> AWS Domain Name Query [Route 53 Resolver]
3. Private hosted zone - Route 53 Resolver
upvoted 1 times

to_to 8 months, 2 weeks ago
Route 53 Resolver : inbound
upvoted 1 times

to_to 8 months, 2 weeks ago
When I organized it slowly, I decided that it was "A" because it was attributed to an account, not a VPC.
upvoted 1 times

veds85 9 months ago
Selected Answer: A
"All VPCs and only need inbound Resolver"
upvoted 1 times

310e976 9 months ago
Answer is A: Please see link below for the solution:
https://docs.aws.amazon.com/whitepapers/latest/hybrid-cloud-dns-options-for-vpc/route-53-resolver-endpoints-and-forwarding-rules.html
upvoted 1 times

masetromain 9 months, 1 week ago
Selected Answer: A
The correct option would be option A:

Associate the private hosted zone to all the VPCs.
Create a Route 53 inbound resolver in the shared services VPC.
Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.
This option will allow the on-premises systems to resolve and connect to cloud.example.com by forwarding the DNS queries to the inbound resolver in the shared services VPC, which will then forward the queries to the private hosted zone. All VPCs will be able to resolve cloud.example.com by resolving the queries through the private hosted zone associated to all VPCs. Additionally, this option takes advantage of the already existing AWS Direct Connect connection between the on-premises corporate network and AWS Transit Gateway, which will provide the highest performance.
upvoted 1 times

c73bf38 9 months, 1 week ago
Selected Answer: A
The best architecture to meet the given requirements with the HIGHEST performance would be Option A:

A. Associate the private hosted zone to all the VPCs. Create a Route 53 inbound resolver in the shared services VPC. Attach all VPCs to the transit gateway and create forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver.

This architecture ensures that all VPCs can resolve the cloud.example.com domain using the private hosted zone. Additionally, it creates a Route 53 inbound resolver in the shared services VPC that can handle DNS resolution requests from on-premises systems through the transit gateway. This setup allows for fast and efficient DNS resolution with minimal latency.
upvoted 1 times

Question #2 *Topic 1*

A company is providing weather data over a REST-based API to several customers. The API is hosted by Amazon API Gateway and is integrated with different AWS Lambda functions for each API operation. The company uses Amazon Route 53 for DNS and has created a resource record of weather.example.com. The company stores data for the API in Amazon DynamoDB tables. The company needs a solution that will give the API the ability to fail over to a different AWS Region.
Which solution will meet these requirements?

A. Deploy a new set of Lambda functions in a new Region. Update the API Gateway API to use an edge-optimized API endpoint with Lambda functions from both Regions as targets. Convert the DynamoDB tables to global tables.

B. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.

C. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a failover record. Enable target health monitoring. Convert the DynamoDB tables to global tables.

D. Deploy a new API Gateway API in a new Region. Change the Lambda functions to global functions. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.

Suggested Answer: *C*

Community vote distribution

robertohyena Highly Voted 2 years, 6 months ago

C.
https://docs.aws.amazon.com/apigateway/latest/developerguide/dns-failover.html
upvoted 16 times

leehjworking 2 years, 2 months ago

Step1 - set up resources - Route 53 failover DNS records for the domain names
upvoted 4 times

c73bf38 Highly Voted 9 months, 1 week ago

The best solution to give the API the ability to fail over to a different AWS Region would be option C:

C. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a failover record. Enable target health monitoring. Convert the DynamoDB tables to global tables.

This solution involves deploying a new API Gateway API and Lambda functions in another region. The company should also convert the DynamoDB tables to global tables to enable cross-region replication of the data. Then, the company should change the Route 53 DNS record to a failover record and enable target health monitoring to automatically route traffic to the new region in the event of a failure or outage in the primary region.
upvoted 9 times

teeee123 Most Recent 4 months, 1 week ago

Selected Answer: C

I think C
upvoted 1 times

TariqKipkemei 8 months ago

Selected Answer: C

Failover routing policy – Use when you want to configure active-passive failover.
upvoted 1 times

masetromain 9 months, 1 week ago

Selected Answer: C

The solution that will meet these requirements is option C:

Deploy a new API Gateway API and Lambda functions in another Region.
Change the Route 53 DNS record to a failover record.
Enable target health monitoring.

Convert the DynamoDB tables to global tables.

This solution will allow the API to failover to a different region, by using Route 53 failover record. The failover record will direct traffic to the primary API endpoint (the one in the primary region) as long as it is healthy. If the primary endpoint becomes unavailable, traffic will be directed to the secondary endpoint (the one in the secondary region). Additionally, by converting the DynamoDB tables to global tables, the data will be available in both regions, which is required for the failover scenario. Target health monitoring can be used to monitor the health of the API Gateway, and when it is determined that the primary endpoint is unavailable, the traffic will be directed to the secondary endpoint.

upvoted 3 times

Sarutobi 9 months, 1 week ago

Selected Answer: C

I also agree with C. But not sure why not B, B is actually pretty good option. No, that I have experience in this specific case; what I normally see is Active/Standby. But option B sounds good because, in theory, we need to have both regions running the current code (Lambda) and if an outage happens we are sure both work, and we don't have stale config/code in the failover region. Sometimes multi-answer does not return the best endpoint for the use case, so that could be something against this solution.

upvoted 3 times

princajen 5 months, 2 weeks ago

Multivalue is used for load balancing, not failover.

upvoted 1 times

edder 9 months, 1 week ago

Selected Answer: B

The answer is B.

A: There is no Route 53, so it cannot be switched in the event of a failure.

C: It's good to change to a failover record, but compared to other questions, there is no step to add a DNS record answer, so you can't switch to a new region.

D: The global function is meaningless.

B: A health check is additionally set, and failover is possible because the corresponding records are not returned in the event of a region failure.

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-configuring.html

upvoted 1 times

ninomfr64 9 months, 1 week ago

Selected Answer: C

Not A. "edge-optimized API endpoint" make use of CloudFront to optimize global each, however API Gateway instance is deployed in a single region thus no ability to fail over to a different AWS Region

Not B. "Route 53 DNS record to a multivalue" implements a active-active scenario, while we are requested to have fail over

Not D. I am not aware of "global function" also "Route 53 DNS record to a multivalue" is not the best fit (see above)

Thus C. is correct has it come with all the required pieces

upvoted 3 times

atirado 9 months, 1 week ago

Selected Answer: C

Option A - Does not provide a way to fail over to a new region but rather a way for API gateway to respond from the region closest to the client

Option B - Does not provide a way to fail over to a new region because when the main region is healthy name resolution will provide 2 possible regions to connect to

Option C - Provides a way to fail over to a new region through the use of a Route 53 failover record and health monitoring and deployment in another region

Option D - Does not provide a way to fail over to a new region because when the main region is healthy name resolution will provide 2 possible regions to connect to

upvoted 5 times

higashikumi 9 months, 1 week ago

Selected Answer: C

To achieve automatic failover for the weather API, the company should deploy a duplicate API Gateway and Lambda functions in a secondary AWS region, then configure a Route 53 failover record that points to both endpoints. This failover record, combined with health checks, will automatically redirect traffic to the secondary region if the primary one fails. Additionally, converting DynamoDB tables to global tables ensures data availability in both regions, allowing the secondary API to function seamlessly during a failover.

upvoted 3 times

amministrazione 10 months ago

D. Deploy a new API Gateway API in a new Region. Change the Lambda functions to global functions. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.

upvoted 1 times

Helpnosense 1 year ago

Selected Answer: D

The changes of A and C are too much, breaking the original security design.

B is wrong because answer B doesn't mention deny SCP on root level is changed. Allow on OU will not win because when allow and deny the same service, explicit deny always wins for the sake of security concerns.

upvoted 1 times

lighthouse85 1 year ago

Selected Answer: C

C, failover health

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, failover record, this is the typical failover configuration on route53. Be careful, chatgpt suggests the option B "multivalue answer"

upvoted 1 times

MoT0ne 1 year, 3 months ago

Selected Answer: C

Choosing C cause you want the API GW and Lambda functions work as a combination behind the DNS with failover, can think of Route53 here as a CDN provider like Cloudflare

upvoted 3 times

abeb 1 year, 7 months ago

C is correct

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: C

failover is required

upvoted 1 times

Question #3 *Topic 1*

A company uses AWS Organizations with a single OU named Production to manage multiple accounts. All accounts are members of the Production OU. Administrators use deny list SCPs in the root of the organization to manage access to restricted services.
The company recently acquired a new business unit and invited the new unit's existing AWS account to the organization. Once onboarded, the administrators of the new business unit discovered that they are not able to update existing AWS Config rules to meet the company's policies.
Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

A. Remove the organization's root SCPs that limit access to AWS Config. Create AWS Service Catalog products for the company's standard AWS Config rules and deploy them throughout the organization, including the new account.

B. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the new account to the Production OU when adjustments to AWS Config are complete.

C. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only. Temporarily apply an SCP to the organization's root that allows AWS Config actions for principals only in the new account.

D. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the organization's root SCP to the Production OU. Move the new account to the Production OU when adjustments to AWS Config are complete.

Suggested Answer: *B*

Community vote distribution

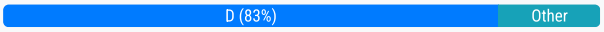

Snip Highly Voted 2 years, 6 months ago

Right answer is D.
An SCP at a lower level can't add a permission after it is blocked by an SCP at a higher level. SCPs can only filter; they never add permissions.
SO you need to create a new OU for the new account assign an SCP, and move the root SCP to Production OU. Then move the new account to production OU when AWS config is done.

upvoted 50 times

robertohyena Highly Voted 2 years, 6 months ago

Answer: D.

Not A: too much overhead and maintenance.
Not B: SCP at Root will still deny Config to the temporary OU.
Not C: Too much overhead to create allow list.

upvoted 20 times

MasterVivek Most Recent 4 days, 7 hours ago

Selected Answer: B

Explanation:
The company uses deny list SCPs at the root level, which means all accounts, including newly added ones, are subject to those restrictions.
The new account needs temporary access to AWS Config actions to update rules, but the current SCPs are blocking that.
Option B allows for a temporary and isolated change:
By creating a separate Onboarding OU, you can apply a more permissive SCP just for that OU.
Once the necessary AWS Config changes are made, the account can be moved to the Production OU, where the standard restrictions apply.
This approach avoids modifying root-level SCPs or converting the entire policy model, which would introduce long-term maintenance overhead.

upvoted 1 times

kvin97 2 weeks, 2 days ago

Selected Answer: D

Answer: D
But it needs to say deny list policy is moved back to the root level, after AWS config process.

upvoted 1 times

calcinator423 1 month, 1 week ago

Selected Answer: B

Here's my take, and let me know if its wrong.

D is wrong because it suggests moving the Root SCP. This is extremely dangerous and generally not advisable, but it also implies that the Root SCP disallows AWS Config changes, which would imply previous administrators also could not change AWS Config. This is an assumption, so I would go with B.

upvoted 1 times

MarkM1 1 month, 1 week ago

Selected Answer: B

B

A: Removing root-level SCPs would loosen security controls across all accounts and increase the risk of misconfiguration. Also, using Service Catalog doesn't solve the issue of updating existing AWS Config rules.

C: Converting from deny list to allow list SCPs across the whole organization is a major change, hard to manage, and could unintentionally block many services. It's not recommended unless you're planning a full org-wide shift.

D: Moving the root SCP to the Production OU changes the scope of restrictions — the root SCP applies to all accounts by default. Moving it to a child OU limits its coverage, which introduces long-term security and management risks.

upvoted 1 times

Cpso 6 months, 1 week ago

Selected Answer: C

the question test knowledge about allow /denied inherit. all ou under root with 'deny' can't allow. So B is fail. C , D is correct but D is less operation.

upvoted 1 times

hspc_ 6 months, 3 weeks ago

Selected Answer: D

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).
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html

upvoted 1 times

masetromain 9 months, 1 week ago

Yes, in option D, the solution is to create a temporary OU named Onboarding for the new account. By creating a new OU for the new account, it allows for a new set of permissions and policies to be applied to this account, separate from the existing Production OU.

Once the new OU is created, an SCP is applied to it to allow AWS Config actions. This SCP allows the new account to make necessary adjustments to AWS Config without being blocked by the existing policies at the root level of the organization.

Then, the root SCP that is blocking these actions is moved to the Production OU, where it will continue to block these actions for all other accounts that are members of the Production OU.

Finally, once the necessary adjustments are made, the new account can be moved to the Production OU, where it will be subject to the existing policies and restrictions.

upvoted 4 times

masetromain 2 years, 5 months ago

This approach is the correct solution because it allows the new account to make necessary adjustments to AWS Config while still adhering to the company's policies, and it does not introduce additional long-term maintenance. The new account will be only in the new OU temporarily, and the SCP blocking AWS Config actions will only be in the root temporarily.

upvoted 1 times

c73bf38 9 months, 1 week ago

The best option to allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance would be option D:

D. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboarding OU to allow AWS Config actions. Move the organization's root SCP to the Production OU. Move the new account to the Production OU when adjustments to AWS Config are complete.

This solution involves creating a temporary OU named Onboarding for the new account and applying an SCP to the Onboarding OU that allows AWS Config actions. The organization's root SCP should be moved to the Production OU, and the new account should be moved to the Production OU when

the adjustments to AWS Config are complete. This approach allows the administrators of the new account to make changes to AWS Config rules while maintaining the current policies in the Production OU.

upvoted 1 times

Ajani 9 months, 1 week ago

Please note Question Constraint: Which option will allow administrators to make changes and continue to enforce the current policies without introducing additional long-term maintenance?

Strategies for using SCPs

You can configure the service control policies (SCPs) in your organization to work as either of the following:

A deny list – actions are allowed by default, and you specify what services and actions are prohibited

An allow list – actions are prohibited by default, and you specify what services and actions are allowed.

upvoted 1 times

sebnzogang 9 months, 1 week ago

Selected Answer: B

D: is not correct, because removing the root SCPs on the production OU means removing all the security rules on the services preventing changes, including changes to the AWS Config rules. and depending on the scenario this will be a security hole for production.

Don't forget that the aim is to introduce the new AWS account into the Production OU with the same configurations and restrictions as the accounts that are already there.

So thanks to the temporary OU on which we have an SCP that authorises actions on AWS Config, we just need to modify the configuration of the new account so that it matches the production requirements. Once the configuration requirements have been met, we move the new account into the production OU.

upvoted 6 times

victorHugo 1 year, 10 months ago

" All accounts are members of the Production OU", therefore we don't need the SCP in root.

upvoted 5 times

dimitry_khan_arc 9 months, 1 week ago

Selected Answer: D

Chosen D.

B is not correct because root having explicit deny will override any explicit allow in its child OU even if allowance is given. Unless I keep Onboarding account under a parent where there is not explicit deny for Config service, Onboarding account can not configure. So, need to move the explicit deny from root account to production account and then keep onboarding account under root.

upvoted 3 times

Dgix 9 months, 1 week ago

This question is ambiguous. If D was formulated like this:

"D. Create a temporary OU named Onboarding for the new account. Apply a Config non-blocking SCP to the Onboarding OU to allow AWS Config actions. Apply the organization's root SCP to the Production OU instead of to the root OU. Move the new account to the Production OU when adjustments to AWS Config are complete."

Then D would be a viable option. However, it isn't, and even if it were, it fails to mention the crucial fact that the Root OU always must have an SCP, which in this case must Allow everything. For someone with some experience this is a given, but as it isn't mentioned, I'd go for B.

However, AWS should reformulate the question and the answers. They are really subpar.

upvoted 2 times

JOKERO 1 year, 3 months ago

AWS Config will still be restricted despite the Allow SCP in Onboarding because of the Deny SCP in the root of the organization

upvoted 3 times

fartosh 1 year, 2 months ago

This sentence:

"Apply the organization's root SCP to the Production OU instead of to the root OU."

solves the issue you mentioned. You can safely move this SCP as the question states that all AWS accounts are in Production OU.

upvoted 1 times

awsylum 9 months, 1 week ago

I don't like any of the answers to be honest. Let's look at D since that's the one most people think is right. The problem with D is that you can't detach the last SCP associated with a root container, OU, or account. There has to be at least one. So, removing the SCP from the root and moving it down to the Production OU is a no-go unless you add a permissable SCP to the root. Check the section on detaching here: https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_attach.html

The only way B is correct is if the reason the new admins don't have access to Config is not because Config is in the Deny List, but because the management account doesn't have the appropriate IAM Policy giving PERMISSION to Config. You need both an IAM Policy and a permissable SCP to have permission and access to a service. But, why wasn't IAM Policy mentioned in choice B. Clearly, without that information, choice B also is not right.

upvoted 1 times

awsylum 1 year, 4 months ago

Also, even if you could remove a root SCP, you would never do that in production. You would never just flat remove an SCP with a Deny list just to give one account access to some service. Even if it's temporary, that's a fatal mistake as the other accounts will not be restricted from certain services they shouldn't have access to.

upvoted 1 times

soulation 4 months, 1 week ago

Before Onboarding OU, this organization only had 1 OU (Production). So moving the SCP from root to Production OU doesn't affect other existing accounts.

upvoted 1 times

awsylum 1 year, 4 months ago

The question mentioned a Deny List architecture, but it didn't specifically say Config was in the Deny List. We are assuming that, which could lead to the wrong answer. Unfortunately, I'm not satisfied with any of the answers. Hopefully, this is a question that would be thrown out from the exam. LOL.

upvoted 1 times

ninomfr64 9 months, 1 week ago

Selected Answer: D

This was not easy for me due to wording, however here is my take:

Not A. here we permanently remove SCPs that limit access to AWS Config, while we are requested to continue to enforce the current policies
Not B. temporary OU and related SCP that allows AWS Config are nested under root where SCPs that limit access to AWS Config are applied. As SCP can only remove permission and not add, this will not work
Not C. converting deny list into allow list here is not beneficial also temporarily apply SCP allowing AWS Config does not meet the request to avoid additional long-term maintenance.

Thus D does the job.

upvoted 2 times

atirado 9 months, 1 week ago

Selected Answer: C

Option A - This option actually rolls out AWS Config across the company which is exactly the opposite of what they are doing

Option B - This option does not work because AWS Config will still be restricted despite the Allow SCP in Onboarding because of the Deny SCP in the root of the organization

Option C - This option allows access to AWS Config in the new business unit and restricts access to everything else. However, the SCP will require regular updates to add new AWS services

Option D - This option applies the correct level of access to each OU without needing updates: Onboarding gets access to AWS Config, Production does not and FullAWSAccess is established at the root after the company's Deny SCP is moved.

upvoted 1 times

Question #4 *Topic 1*

A company is running a two-tier web-based application in an on-premises data center. The application layer consists of a single server running a stateful application. The application connects to a PostgreSQL database running on a separate server. The application's user base is expected to grow significantly, so the company is migrating the application and database to AWS. The solution will use Amazon Aurora PostgreSQL, Amazon EC2 Auto Scaling, and Elastic Load Balancing.
Which solution will provide a consistent user experience that will allow the application and database tiers to scale?

A. Enable Aurora Auto Scaling for Aurora Replicas. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

B. Enable Aurora Auto Scaling for Aurora writers. Use an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled.

C. Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.

D. Enable Aurora Scaling for Aurora writers. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

Suggested Answer: *C*

Community vote distribution

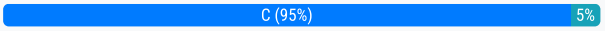

robertohyena Highly Voted 2 years, 6 months ago

C.
- Aurora writers is a distractor.
- Single master mode only has read replica - with Aurora replicas.
- Multi master mode, not in the options
- NLB does not support round robin and least outstanding algorithm

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Integrating.AutoScaling.html
upvoted 29 times

c73bf38 Highly Voted 9 months, 1 week ago

Selected Answer: C

The best solution to provide a consistent user experience that will allow the application and database tiers to scale would be option C:

C. Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.

This solution involves enabling Aurora Auto Scaling for Aurora Replicas to automatically add and remove read replicas to match the application's workload. The solution also uses an Application Load Balancer to distribute traffic to the application layer, with the round robin routing algorithm to balance the traffic evenly across multiple instances. Sticky sessions should be enabled to maintain session affinity for each user, allowing for a consistent user experience.
upvoted 18 times

TariqKipkemei Most Recent 8 months ago

Selected Answer: C

Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky sessions enabled
upvoted 2 times

masetromain 9 months, 1 week ago

Selected Answer: C

C is correct. This solution will provide a consistent user experience by using an Application Load Balancer with the round robin routing algorithm and sticky sessions enabled. This allows the application and database tiers to scale by using Aurora Auto Scaling for Aurora Replicas. This will ensure that the application is able to handle the increased user base while maintaining a consistent user experience. The use of an Application Load Balancer also allows for better routing of traffic to the available Aurora Replicas.
upvoted 2 times

ninomfr64 9 months, 1 week ago

Selected Answer: C

- Auto Scaling for Aurora writers does not exists (distractor)
- NLB does not support least outstanding requests routing algorithm (it only supports Flow Hash)
- NLB does not allow to enable Sticky Sessions, this is always enabled with Flow Hash where each TCP/UDP connection is routed to a single target for the life of the connection

Thus C is correct

upvoted 3 times

atirado 9 months, 1 week ago

Selected Answer: C

Option A - Allows the tiers to grow but NLB does not make load balancing decisions that way

Option B - No such thing as Aurora Autoscaling for Aurora Writers

Option C - Allows the tiers to grow and ALB using sticky sessions provides consistent user experience

Option D - No such thing as Aurora Autoscaling for Aurora Writers

Note: The application is web-based so choosing ALB shouldn't be an issue.

upvoted 3 times

amministrazione 10 months ago

D. Enable Aurora Scaling for Aurora writers. Use a Network Load Balancer with the least outstanding requests routing algorithm and sticky sessions enabled.

upvoted 1 times

Bereket 1 year ago

Selected Answer: C

C, Enable Aurora Auto Scaling for Aurora Replicas

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, Enable Aurora Auto Scaling for Aurora Replicas

upvoted 1 times

MoT0ne 1 year, 3 months ago

Selected Answer: C

Single writer: In an Aurora PostgreSQL DB cluster, there is only one writer instance at a time. All write operations, such as INSERT, UPDATE, and DELETE statements, are directed to the writer instance.

upvoted 4 times

GNB2024 1 year, 4 months ago

Selected Answer: C

It's C

upvoted 1 times

liux99 1 year, 6 months ago

B, D are distractor, as there is no writer replica in aurora autoscale.
NLB does not support sticky session so A is out. The anwser is C.

upvoted 2 times

rhinozD 1 year, 4 months ago

NLB Sticky Session: https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html#sticky-sessions

upvoted 2 times

abeb 1 year, 7 months ago

C Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky session

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: C

Aurora - AS only for read replicas. NLB doesn't support the least outstanding requests or round-robin algorithms, only flow hash is supported.
upvoted 1 times

ansgohar 1 year, 9 months ago

Selected Answer: C

C. Enable Aurora Auto Scaling for Aurora Replicas. Use an Application Load Balancer with the round robin routing and sticky sessions enabled.
upvoted 1 times

rsn 1 year, 9 months ago

Selected Answer: A

NLB scales better than ALB. Also least outstandind requests algorithm works better than round robin algorith. Any thougts?
upvoted 2 times

Ganshank 1 year, 9 months ago

The correct answer is whatever the examiner says it is. Depending on how you look at it either A or C can be the correct answer.
NLB scales better and supports LOR algorithm which are both factors in its favor, however stickiness is not supported for TLS connections in NLBs. While this has not been called out explicitly, I doubt anyone in today's world would support non-TLS connections to their applications. If that turns out to be a dealbreaker, then the only option is C, to use ALB, however round-robin doesn't guarantee the best performance especially where stickiness is concerned.
Your call.
upvoted 3 times

dimitry_khan_arc 1 year, 10 months ago

Selected Answer: C

write replica is distractor. NLB does not support round robin
upvoted 2 times

Question #5 *Topic 1*

A company uses a service to collect metadata from applications that the company hosts on premises. Consumer devices such as TVs and internet radios access the applications. Many older devices do not support certain HTTP headers and exhibit errors when these headers are present in responses. The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices, which the company identified by the User-Agent headers.
The company wants to migrate the service to AWS, adopt serverless technologies, and retain the ability to support the older devices. The company has already migrated the applications into a set of AWS Lambda functions.
Which solution will meet these requirements?

A. Create an Amazon CloudFront distribution for the metadata service. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB. Configure the ALB to invoke the correct Lambda function for each type of request. Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header.

B. Create an Amazon API Gateway REST API for the metadata service. Configure API Gateway to invoke the correct Lambda function for each type of request. Modify the default gateway responses to remove the problematic headers based on the value of the User-Agent header.

C. Create an Amazon API Gateway HTTP API for the metadata service. Configure API Gateway to invoke the correct Lambda function for each type of request. Create a response mapping template to remove the problematic headers based on the value of the User-Agent. Associate the response data mapping with the HTTP API.

D. Create an Amazon CloudFront distribution for the metadata service. Create an Application Load Balancer (ALB). Configure the CloudFront distribution to forward requests to the ALB. Configure the ALB to invoke the correct Lambda function for each type of request. Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of the User-Agent header.

Suggested Answer: *B*

Community vote distribution

EricZhang Highly Voted 2 years, 6 months ago

A. The only difference between A and D is CloudFront function vs Lambda@Edge. In this case the CloudFront function can remove the response header based on request header and much faster/light-weight.

upvoted 64 times

RyGuy2025 4 months, 4 weeks ago

If you read the solution - it does not reference a CloudFront Function, it references a Cloud Front Distribution, which is not the same. That is why B is the best answer.

upvoted 1 times

RyGuy2025 4 months, 4 weeks ago

Where the function is integrated - it is already past the ALB.

upvoted 1 times

vn_thanhtung 1 year, 10 months ago

After read, answer A "Create a CloudFront function to remove the problematic headers based on the value of the User-Agent header" not really clear and fuzzy, "The company has configured an on-premises load balancer to remove the unsupported headers from responses sent to older devices" => "Create a Lambda@Edge function that will remove the problematic headers in response to viewer requests based on the value of the User-Agent header" => D make sence

upvoted 13 times

masetromain Highly Voted 2 years, 6 months ago

I think this is answer D: Lambda@Edge can modify headers
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html

upvoted 30 times

vn_thanhtung 1 year, 10 months ago

Agree D

upvoted 5 times

ninomfr64 1 year, 6 months ago

Agree on D, but also CloudFront Function can manipulate headers
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cloudfront-functions.html#:~:text=cache%20hit%20ratio.-,Header%20manipulation,-%E2%80%93%20You%20can%20insert
upvoted 3 times

0dc6cac Most Recent 1 week, 1 day ago

Selected Answer: B

There is no way A or D are the answers here....ALB + lambda target groups are NOT SERVERLESS...also we don't know the exact nature of the application, it might not be efficient or even possible to use ALB + lambda, if it involves API keys or some complex routing mechanism. AWS would NEVER prefer that we use ALB + lambda TGs to do something like this.

I'd say B is most likely the answer.
upvoted 1 times

jimee11 1 month, 3 weeks ago

Selected Answer: A

CloudFreont functions provide better solution approach for lightweight tasks i.e. header manipulation and basic URL rewrites.
upvoted 1 times

Hello43638 2 months, 1 week ago

Selected Answer: D

➡ This is where CloudFront Functions falls short:
While it can manipulate response headers, it cannot access request headers during the response phase.
So it can't conditionally remove headers from responses based on User-Agent.
upvoted 1 times

unbornfroyo 2 months, 1 week ago

Selected Answer: D

TL;DR:
You need to remove specific headers based on User-Agent in the response, best done at the edge.

Lambda@Edge with CloudFront is the only option here that fully supports this requirement.
upvoted 1 times

chucky41_1 2 months, 2 weeks ago

Selected Answer: C

Answer: C
Question mentions to use serverless technologies,

API Gateway and Lambda considerd as serverless. so answer A & D can be ommited and B or C are the remaining options, but since Option C mentions mapping templates, which has ability to modify the request. that is the probable answer.
upvoted 1 times

cloudlab 3 months ago

Selected Answer: A

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/edge-functions-choosing.html
upvoted 1 times

ParamD 3 months, 2 weeks ago

Selected Answer: A

CloudFront functions are the easiest way to manipulate response headers
upvoted 1 times

BennyMao 3 months, 3 weeks ago

Selected Answer: D

CloudFront functions cannot modify response headers, which is a key requirement.
upvoted 1 times

soulation 4 months, 1 week ago

Selected Answer: A

While D might work, it's overkill for this use case. Cloudfront function is 600% cheaper, less latency, less code size. Correct answer is A.
upvoted 1 times

RyGuy2025 4 months, 4 weeks ago

Selected Answer: B

I'm going to wade in here and vote B as the most efficient and cost effective deployment option - here's why:
REST APIs natively support header modifications and API Gateway responses can be customized based on conditions (major feature).
You can easily integrate the already moved Lamda functions - no additional components required.
Lastly, the solution meets the serverless requirement and is more cost effective than the others.

upvoted 2 times

RyGuy2025 4 months, 4 weeks ago

Following this up to say after further research I believe B is wrong - no way to remove though.. sry ppl - API Gateway REST APIs cannot natively modify response headers based on the User-Agent header in a conditional way - sry ppl

upvoted 1 times

pk0619 6 months, 2 weeks ago

Selected Answer: A

CloudFront Functions is ideal for lightweight, short-running functions for the following use cases:

Header manipulation – Insert, modify, or delete HTTP headers in the request or response. For example, you can add a True-Client-IP header to every request

upvoted 2 times

julianocaldeira 7 months, 2 weeks ago

The correct answer is D.

Here's why:

Requirement: The solution must:

Use serverless technologies.
Retain the ability to support older devices by removing problematic headers based on the User-Agent header.
Analysis of the options:

A. CloudFront with an ALB and a CloudFront function:

CloudFront functions can modify headers, but they are limited to request handling and cannot modify response headers.
This does not fully meet the requirement to remove response headers based on the User-Agent header.

upvoted 1 times

Kirkster 5 months, 3 weeks ago

According to the docs, CloudFront functions can, "Insert, modify, or delete HTTP headers in the request or response." So A is the better choice.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/edge-functions-choosing.html

upvoted 2 times

ad11934 7 months, 2 weeks ago

Option A as per https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/edge-functions-choosing.html -- seems lambda@edge is not needed as cloudfront function can be quick for modifying the response headers and handles higher request rates

upvoted 1 times

CGarces 7 months, 2 weeks ago

Selected Answer: A

It's A.
Cloudfront funtions are faster and cheaper than Lambda@Edge.
The use case in the question (add headers if missing from old devices) is documented on the AWS samples
https://github.com/aws-samples/amazon-cloudfront-functions/tree/main/add-origin-header

upvoted 3 times

MChal 7 months, 3 weeks ago

Option A or D are viable. I went with D because Lambda@Edge offers better performance since it runs closer to the end-users, reducing latency. Lambda@Edge can also handle more complex logic and is always more suitable for modifying headers based on User-Agent. While A could work, D offers a more robust and scalable solution that better aligns with the company's requirements to support older devices and leverage serverless technologies effectively.

upvoted 1 times

Question #6 *Topic 1*

A retail company needs to provide a series of data files to another company, which is its business partner. These files are saved in an Amazon S3 bucket under Account A, which belongs to the retail company. The business partner company wants one of its IAM users, User_DataProcessor, to access the files from its own AWS account (Account B).
Which combination of steps must the companies take so that User_DataProcessor can access the S3 bucket successfully? (Choose two.)

A. Turn on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account A.

B. In Account A, set the S3 bucket policy to the following:

```
{
    "Effect": "Allow",
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

C. In Account A, set the S3 bucket policy to the following:

```
{
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
    },
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::AccountABucketName/*"
    ]
}
```

D. In Account B, set the permissions of User_DataProcessor to the following:

```
{
    "Effect": "Allow",
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": "arn:aws:s3:::AccountABucketName/*"
}
```

E. In Account B, set the permissions of User_DataProcessor to the following:

```
{
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::AccountB:user/User_DataProcessor"
    },
    "Action": [
        "s3:GetObject",
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::AccountABucketName/*"
    ]
}
```

Suggested Answer: *D*

Community vote distribution

C (63%) D (34%)

robertohyena Highly Voted 9 months, 1 week ago

Answer: C & D

Source:
https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/

https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-walkthroughs-managing-access-example4.html
upvoted 34 times

higashikumi Highly Voted 9 months, 1 week ago
C & D

To allow User_DataProcessor to access the S3 bucket from Account B, the following steps need to be taken:

In Account A, set the S3 bucket policy to allow access to the bucket from the IAM user in Account B. This is done by adding a statement to the bucket policy that allows the IAM user in Account B to perform the necessary actions (GetObject and ListBucket) on the bucket and its contents.

In Account B, create an IAM policy that allows the IAM user (User_DataProcessor) to perform the necessary actions (GetObject and ListBucket) on the S3 bucket and its contents. The policy should reference the ARN of the S3 bucket and the actions that the user is allowed to perform.

Note: turning on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account A is not necessary for this scenario as it is typically used for allowing web browsers to access resources from different domains.
upvoted 19 times

anhdadenpp Most Recent 2 months, 1 week ago
Selected Answer: D
Answer : CD
Detail : https://docs.aws.amazon.com/res/latest/ug/S3-buckets-cross-account-access.html
upvoted 1 times

DhirajBansal 7 months, 1 week ago
Selected Answer: C
This policy will give access to User in B account for Bucket in A account.
upvoted 1 times

Jorkaef 7 months, 2 weeks ago
The correct combination of steps for this scenario are:

C. In Account A, set the S3 bucket policy to the following:

E. In Account B, set the permissions of User_DataProcessor to the following:

Here's why these are the correct steps:

Step C: The bucket policy in Account A (the retail company) needs to explicitly allow access to the IAM user from Account B (the business partner). This policy grants the necessary permissions to User_DataProcessor from Account B to access the S3 bucket in Account A.

Step E: In Account B (the business partner's account), the IAM user User_DataProcessor needs to be granted permissions to access S3 resources. This IAM policy allows the user to perform the necessary S3 actions.
upvoted 3 times

TariqKipkemei 8 months ago
Selected Answer: C
C & D.
In Account A, set the S3 bucket policy to allow only 'User_DataProcessor' from Account B access.
In Account B, set the permissions of User_DataProcessor to allow access to S3 bucket in Account A.
upvoted 1 times

85b5b55 8 months, 1 week ago
Answer: C & D
upvoted 1 times

atirado 9 months, 1 week ago

Selected Answer: C

Option A - CORS does not address cross-account access to S3 buckets

Option B - This option would not work because the bucket policy is missing the Principal

Option C - This option provides a valid S3 bucket policy that grants access to User_DataProcessor

Option D - These permissions allow User_DataProcessor to get objects out of the bucket

Option E - This option would not work because it is not a valid IAM policy

upvoted 1 times

amministrazione 10 months ago

C. In Account A, set the S3 bucket policy to the following:

D. In Account B, set the permissions of User_DataProcessor to the following:

upvoted 1 times

dEgYnIDA 11 months, 2 weeks ago

Selected Answer: D

The question says Choose two. The answer is C & D.

upvoted 1 times

kpcert 1 year ago

Selected Answer: C

Ans C and D

2 Options have to be selected

upvoted 1 times

kpcert 1 year ago

Ans - C and D

2 Options have to be selected

upvoted 1 times

MoT0ne 1 year, 3 months ago

Selected Answer: C

Cross-Origin Resource Sharing (CORS) is a security feature in Amazon S3 that allows you to control access to your S3 resources from a different domain (origin) than the one serving the resources. CORS defines a way for client web applications running in one origin to interact with resources in a different origin, which is otherwise restricted by the same-origin policy enforced by web browsers.

upvoted 1 times

Dgix 1 year, 4 months ago

C and D.

upvoted 1 times

awsylum 1 year, 4 months ago

The answer is C and D. You need to give the IAM User in Account B an IAM Policy and you need to give a Bucket Policy in Account A.

Who is maintaining this database of questions? Someone needs to seriously set the correct answers before making a lot of people confused and potentially screw up their exam.

upvoted 2 times

chelbsik 1 year, 4 months ago

Selected Answer: D

Correct answer: C and D

Adding my vote for D to balance the result

Moderator, please fix the vote in this ticket.

upvoted 2 times

ftaws 1 year, 4 months ago

why we need two steps? I think that we get only one from resource-based policy or identity-based policy.

upvoted 1 times

Question #7 *Topic 1*

A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.
Which solution will meet these requirements MOST cost-effectively?

A. Upload the container images to AWS Lambda as functions. Configure a concurrency limit for the associated Lambda functions to handle the expected peak load. Configure two separate Lambda integrations within Amazon API Gateway: one for production and one for testing.

B. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.

C. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Kubernetes Service (Amazon EKS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the EKS clusters.

D. Upload the container images to AWS Elastic Beanstalk. In Elastic Beanstalk, create separate environments and deployments for production and testing. Configure two separate Application Load Balancers to direct traffic to the Elastic Beanstalk deployments.

Suggested Answer: *B*

Community vote distribution

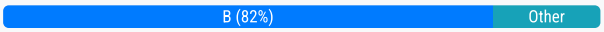

masetromain Highly Voted 9 months, 1 week ago

Selected Answer: B

B. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.
This option meets the requirement of using a serverless architecture by utilizing the Fargate launch type for the ECS clusters, which allows for automatic scaling of the containers based on the expected load. It also allows for separate deployments for production and testing by configuring separate ECS clusters and Application Load Balancers for each environment. This option also minimizes operational complexity by utilizing ECS and Fargate for the container orchestration and scaling.
upvoted 22 times

zhangyu20000 Highly Voted 2 years, 6 months ago

Answer is A. ABC all works but A is most COST EFFECTIVE
upvoted 15 times

masetromain 2 years, 6 months ago

Is true but " you can now package and deploy Lambda functions as container images of up to 10 GB in size." the size is not specified, personally I find it too small
upvoted 3 times

Kirkster 5 months, 2 weeks ago

10 GB is ENORMOUS for a container image. Even most Windows Server OS container images (which won't work in Lambda anyway) are smaller than that, and it's very very rare to see a Linux container image over ~1 GB (typically they are a few hundred MB)
upvoted 1 times

anita_student 2 years, 4 months ago

10GB image is too small for what? I'm curious how do you containerise those images?
I'd say the average image size is ~300-400MB
upvoted 5 times

zhangyu20000 2 years, 6 months ago

https://aws.amazon.com/blogs/aws/new-for-aws-lambda-container-image-support/
upvoted 3 times

anita_student 2 years, 4 months ago

Yes, would be cheap, but can't run a web app from Lambda
upvoted 5 times

Kirkster 5 months, 2 weeks ago
Of course you can run web applications from Lambda, with API Gateway (or Lambda HTTP URL) in front of it. You have to refactor a little, unless you're using ASP.NET Core, which can run nearly unmodified in Lambda using the Amazon.Lambda.AspNetCoreServer package.
upvoted 1 times

yuyuyuyuyu 2 years, 5 months ago
I do not think A is the right answer.
Because image must be upload to the ECR.
upvoted 4 times

sasivarenan Most Recent 1 week ago
Selected Answer: B
A is definitely no for hosting web applications, web applications typically make multiple HTTP requests where stick session are required which is not supported by API Gateway. So B looks perfect
upvoted 1 times

0dc6cac 2 weeks, 2 days ago
Selected Answer: A
I think it's A, you can upload docker images to lambda, and it's serverless. I don't think ECS can count as serverless in this case.
upvoted 1 times

rhuanca 3 weeks, 3 days ago
Selected Answer: A
looks like now since 2020 lambda support container images
https://aws.amazon.com/blogs/aws/new-for-aws-lambda-container-image-support/?utm_source=chatgpt.com
upvoted 2 times

SBoksh 1 month, 1 week ago
Selected Answer: A
https://docs.aws.amazon.com/lambda/latest/dg/images-create.html
upvoted 1 times

zhen234 4 months, 3 weeks ago
Selected Answer: A
most cost-effective and lest operational overhead
upvoted 1 times

Kirkster 5 months, 2 weeks ago
Selected Answer: B
I was initially torn between A and B, but answer A says to upload the container image to Lambda, which isn't possible - to use a container with Lambda, you still upload the image to ECR. Answer D (Beanstalk) isn't the most cost-effective for running containers.

So between B and C (ECS vs EKS), ECS has less operational overhead, and also doesn't require the master node to be running, which means ECS will likely be very slightly cheaper, and have less operational work.
upvoted 1 times

fbukevin 6 months ago
Selected Answer: B
At the time I consider B & C without doubt. But finally consider to migration efforts, I choose B. I don't really consider the cost between B & C.
upvoted 1 times

attila9778 6 months, 4 weeks ago
Selected Answer: B
AWS Fargate launch type for Amazon ECS is indeed a pay-per-use model. With Fargate, you pay for the amount of vCPU and memory resources that your containerized application requests. The billing is based on the resources used from the time your container images are pulled until the Amazon ECS task terminates, rounded up to the nearest second, with a minimum charge of one minute12.

This model allows you to focus on building and managing your applications without worrying about managing the underlying infrastructure, making it a convenient and scalable option for many use cases.
See: https://aws.amazon.com/ecs/pricing/
upvoted 1 times

TariqKipkemei 8 months ago

Selected Answer: B

"microservices that run on containers, with a serverless architecture that minimizes operational complexity and costs" = Amazon Elastic Container Service (Amazon ECS)

upvoted 1 times

85b5b55 8 months, 1 week ago

Serverless and Most Cost-Effective Solutions - B

upvoted 1 times

AWSum1 8 months, 1 week ago

Selected Answer: A

Choosing A, Fargate had a minimum charge per minute. Lambda is per invocation. So lambda Is probably cheaper in THIS case

upvoted 1 times

AWSum1 8 months, 1 week ago

Adding to this, the questions states minimal operation complexity. ECR + ECS + ELB seems like quite a lot to manage

upvoted 1 times

c73bf38 9 months, 1 week ago

Selected Answer: B

Option B is the most cost-effective solution for the following reasons:

The use of Fargate, a serverless compute engine for containers, eliminates the need for managing and scaling the underlying infrastructure. This minimizes operational complexity and reduces costs as the resources are used only when required.
Auto scaling ensures that the application scales up and down based on the load, providing the required performance and availability without incurring additional costs.
Amazon ECS is a simpler and more cost-effective solution than Amazon EKS, which requires more management and additional resources to operate the Kubernetes control plane.
Using Application Load Balancers to direct traffic to the ECS clusters ensures high availability and fault tolerance.

upvoted 4 times

c73bf38 2 years, 4 months ago

Changing to A, B is not serverless and cost-effective.

upvoted 1 times

bcx 2 years ago

Fargate is serverless by definition.

upvoted 2 times

Jonalb 9 months, 1 week ago

Selected Answer: B

Explanation:

Amazon ECS with Fargate: By uploading the container images to Amazon ECR and using Amazon ECS with the Fargate launch type, you can run the microservices in containers without having to manage the underlying infrastructure. Fargate automatically scales the containers based on the load.
Separate Production and Testing Environments: With two separate auto-scaled Amazon ECS clusters, you can have dedicated environments for production and testing, ensuring isolation and allowing for separate deployments and configurations.
Application Load Balancers (ALB): Configuring two separate ALBs allows you to direct traffic to the appropriate ECS clusters. This ensures proper routing of requests between the production and testing environments.
Option B provides a cost-effective solution by utilizing the serverless nature of Fargate, which eliminates the need to provision and manage EC2 instances explicitly. It also allows for separate environments, easy scalability, and traffic routing using ALBs, providing flexibility and minimizing operational complexity.

upvoted 2 times

atirado 9 months, 1 week ago

Selected Answer: B

Option A - This option might not work. AWS Lambda provides a cheap option to run containers however nothing is said about execution times could be a concern, i.e. AWS Lambda only provides 15 minutes of execution time

Option B - This option will work. ALB, ECR, ECS and Fargate in combination will deliver a running solution.

Option C - This option will work. ALB, ECR, EKS and Fargate will deliver a running solution.

Option D - This option will work: Beanstalk will rely on ECS to run the containers. See https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

Cheapest option is B.

upvoted 6 times

amministrazione 10 months ago

B. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.

upvoted 1 times

Question #8 *Topic 1*

A company has a multi-tier web application that runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB and the Auto Scaling group are replicated in a backup AWS Region. The minimum value and the maximum value for the Auto Scaling group are set to zero. An Amazon RDS Multi-AZ DB instance stores the application's data. The DB instance has a read replica in the backup Region. The application presents an endpoint to end users by using an Amazon Route 53 record.
The company needs to reduce its RTO to less than 15 minutes by giving the application the ability to automatically fail over to the backup Region. The company does not have a large enough budget for an active-active strategy.
What should a solutions architect recommend to meet these requirements?

A. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.

B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.

C. Configure the Auto Scaling group in the backup Region to have the same values as the Auto Scaling group in the primary Region. Reconfigure the application's Route 53 record with a latency-based routing policy that load balances traffic between the two ALBs. Remove the read replica. Replace the read replica with a standalone RDS DB instance. Configure Cross-Region Replication between the RDS DB instances by using snapshots and Amazon S3.

D. Configure an endpoint in AWS Global Accelerator with the two ALBs as equal weighted targets. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Create an Amazon CloudWatch alarm that is based on the HTTPCode_Target_5XX_Count metric for the ALB in the primary Region. Configure the CloudWatch alarm to invoke the Lambda function.

Suggested Answer: *B*

Community vote distribution

B (100%)

masetromain Highly Voted 2 years, 6 months ago

Selected Answer: B

I go with B
https://docs.amazonaws.cn/en_us/Route53/latest/DeveloperGuide/welcome-health-checks.html
upvoted 19 times

masetromain 2 years, 5 months ago

B is correct, because it meets the company's requirements for reducing RTO to less than 15 minutes and not having a large budget for an active-active strategy.

In this solution, the company creates an AWS Lambda function in the backup region which promotes the read replica and modifies the Auto Scaling group values. Route 53 is configured with a health check that monitors the web application and sends an Amazon SNS notification to the Lambda function when the health check status is unhealthy. The Route 53 record is also updated with a failover policy that routes traffic to the ALB in the backup region when a health check failure occurs. This way, when the primary region goes down, the failover policy triggers and traffic is directed to the backup region, ensuring a quick recovery time.
upvoted 17 times

TariqKipkemei Most Recent 8 months ago

Selected Answer: B

Option B is the least costly active-passive strategy
upvoted 1 times

Untamables 9 months, 1 week ago

Selected Answer: B

I Vote B.
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html

Option A, C and D are wrong. The latency-based routing and endopoint weights should be used for active/active strategy.
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy-latency.html
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-endpoints-endpoint-weights.html
upvoted 4 times

higashikumi 9 months, 1 week ago
The best option to meet the requirements and reduce RTO to less than 15 minutes is to choose option B.

Option B involves creating an AWS Lambda function in the backup region to promote the read replica and modify the Auto Scaling group values. Additionally, Route 53 can be configured with a health check that monitors the web application and sends an Amazon SNS notification to the Lambda function when the health check status is unhealthy. The application's Route 53 record can be updated with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.

This option is cost-effective as it does not require an active-active strategy, and it uses AWS services to minimize the RTO. The Lambda function can be invoked to promote the read replica in the backup region, and the Auto Scaling group values can be updated to launch EC2 instances in the backup region. Furthermore, the Route 53 health check feature can be used to monitor the web application and initiate the failover process.
upvoted 1 times

atirado 9 months, 1 week ago
Selected Answer: B
Option A - This option will not work as needed: The client will get errors when the closest region is the application's backup region

Option B - This option implements an active-passive strategy as needed: When the health check fails, Route 53 will resolve to the backup region and the Lambda function will ensure the backup region has resources to function

Option C - This option implements an active-active strategy

Option D - This option will not work as needed: The client will get errors 50% of the time
upvoted 3 times

GreyBox1 9 months, 3 weeks ago
Selected Answer: B
B is right.
upvoted 1 times

amministrazione 10 months ago
B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.
upvoted 1 times

Bereket 1 year ago
Selected Answer: B
Correct answer B
upvoted 1 times

gofavad926 1 year, 3 months ago
Selected Answer: B
B for sure
upvoted 1 times

Vaibs099 1 year, 5 months ago
This explains Lambda promoting backup read replica in other region - https://medium.com/ankercloud-engineering/aws-lambda-promoting-rds-read-replica-on-cross-region-using-aws-lambda-113db758869
upvoted 1 times

ftaws 1 year, 5 months ago
why we need Lambda Function ? Is it enough a Route 53 failover policy ?
upvoted 1 times

rhinozD 1 year, 4 months ago
What about RDS failover?
You need lambda to promote read replica.

upvoted 1 times

ninomfr64 1 year, 6 months ago

Selected Answer: B

The problem is not detecting the right answer, but reading quickly enough trough all the words in the question!

upvoted 1 times

jainparag1 1 year, 7 months ago

Selected Answer: B

B satisfies all the requirements.

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: B

Health check is a metric, hence alarms can be executed, and alarms are integrated with SNS, SNS integrated with lambda. This sounds weird, but it will work.

upvoted 1 times

ansgohar 1 year, 9 months ago

Selected Answer: B

B. Create an AWS Lambda function in the backup Region to promote the read replica and modify the Auto Scaling group values. Configure Route 53 with a health check that monitors the web application and sends an Amazon Simple Notification Service (Amazon SNS) notification to the Lambda function when the health check status is unhealthy. Update the application's Route 53 record with a failover policy that routes traffic to the ALB in the backup Region when a health check failure occurs.

upvoted 2 times

dimitry_khan_arc 1 year, 10 months ago

Selected Answer: B

Health check+SNS. This does not need to have active-active which satisfy the rquirement.

upvoted 1 times

NikkyDicky 2 years ago

it's a B again

upvoted 1 times

Question #9 *Topic 1*

A company is hosting a critical application on a single Amazon EC2 instance. The application uses an Amazon ElastiCache for Redis single-node cluster for an in-memory data store. The application uses an Amazon RDS for MariaDB DB instance for a relational database. For the application to function, each piece of the infrastructure must be healthy and must be in an active state.
A solutions architect needs to improve the application's architecture so that the infrastructure can automatically recover from failure with the least possible downtime.
Which combination of steps will meet these requirements? (Choose three.)

A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.

B. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are configured in unlimited mode.

C. Modify the DB instance to create a read replica in the same Availability Zone. Promote the read replica to be the primary DB instance in failure scenarios.

D. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.

E. Create a replication group for the ElastiCache for Redis cluster. Configure the cluster to use an Auto Scaling group that has a minimum capacity of two instances.

F. Create a replication group for the ElastiCache for Redis cluster. Enable Multi-AZ on the cluster.

Suggested Answer: *ADF*

Community vote distribution

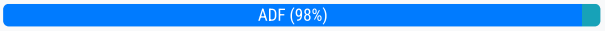

masetromain Highly Voted 2 years, 6 months ago

Selected Answer: ADF

I go with ADF
https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html
upvoted 16 times

masetromain 9 months, 1 week ago

A. Using an Elastic Load Balancer (ELB) to distribute traffic across multiple EC2 instances can help ensure that the application remains available in the event that one of the instances becomes unavailable. By configuring the instances as part of an Auto Scaling group with a minimum capacity of two instances, you can ensure that there is always at least one healthy instance to handle traffic.

D. Modifying the DB instance to create a Multi-AZ deployment that extends across two availability zones can help ensure that the database remains available in the event of a failure. In the event of a failure, traffic will automatically be directed to the secondary availability zone, reducing the amount of downtime.

F. Creating a replication group for the ElastiCache for Redis cluster and enabling Multi-AZ can help ensure that the in-memory data store remains available in the event of a failure. This will allow traffic to be automatically directed to the secondary availability zone, reducing the amount of downtime.
upvoted 11 times

spencer_sharp 2 years, 6 months ago

Why C is wrong?
upvoted 3 times

Karamen 1 year, 10 months ago

let suppose in case one of used AZ is failed?
upvoted 1 times

masetromain 2 years, 5 months ago

Other options like B. and C. does not meet the requirement because the instances are configured in unlimited mode, it will not be possible to ensure that there is always at least one healthy instance to handle traffic if there is a failure.
upvoted 1 times

God_Is_Love 2 years, 4 months ago

Issue with C - Read replica in the same AZ does not sound High availability

upvoted 6 times

dtha1002 2 years, 1 month ago

in question "can automatically recover from failure with the least possible downtime"

C is correct but D is least possible downtime

upvoted 1 times

teeee123 Most Recent 4 months, 1 week ago

Selected Answer: ADF

ADF try it

upvoted 1 times

TariqKipkemei 8 months ago

Selected Answer: ADF

Key words: each piece of the infrastructure must be healthy and must be in an active state, can automatically recover = Auto Scaling group and Multi-AZ deployment

upvoted 1 times

higashikumi 9 months, 1 week ago

A, D, E are the correct options to meet the requirements.

Option A is correct because an Auto Scaling group with a minimum capacity of two instances and an Elastic Load Balancer distributing traffic across them can provide high availability and automatic recovery from failure.

Option D is correct because a Multi-AZ deployment for the RDS instance will ensure that there is a synchronized standby copy of the database in a separate Availability Zone that can be used for automatic failover.

Option E is correct because configuring an Auto Scaling group for the ElastiCache for Redis cluster will ensure that there is at least one available node at all times, and automatic recovery can be achieved by launching new instances to replace any failed nodes.

upvoted 1 times

marszalekm 1 year, 5 months ago

There isn't such a thing like "Auto Scaling group for the ElastiCache for Redis", there is a "Replication Group"

upvoted 2 times

Maja1 9 months, 1 week ago

Selected Answer: ADF

I wasn't sure if E or F was correct until I read this:

"This replacement results in some downtime for the cluster, but if Multi-AZ is enabled, the downtime is minimized. The role of primary node will automatically fail over to one of the read replicas. There is no need to create and provision a new primary node, because ElastiCache will handle this transparently. This failover and replica promotion ensure that you can resume writing to the new primary as soon as promotion is complete."

https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html

upvoted 4 times

atirado 9 months, 1 week ago

Selected Answer: ADF

Option A - Ensures there is always at least a healthy instance responding to requests. Nothing is said about whether the Auto Scaling Group includes multiple AZs (but it must)

Option B - No such thing as EC2 Unlimited Mode

Option C - Does not provide a place to fail over to

Option D - Provides a place to fail over to

Option E - Does not provide a place to fail over to

Option F - Provides a place to fail over to

Choose A, D, F

upvoted 2 times

amministrazione 10 months ago

A. Use an Elastic Load Balancer to distribute traffic across multiple EC2 instances. Ensure that the EC2 instances are part of an Auto Scaling group that has a minimum capacity of two instances.

D. Modify the DB instance to create a Multi-AZ deployment that extends across two Availability Zones.

F. Create a replication group for the ElastiCache for Redis cluster. Enable Multi-AZ on the cluster.

upvoted 1 times

[Removed] 1 year, 1 month ago

ElastiCache for Redis Auto Scaling is limited to the following:

Redis (cluster mode enabled) clusters running Redis engine version 6.0 onwards

E is out

upvoted 1 times

joshnort 1 year, 1 month ago

Selected Answer: ADF

Satisfies the High Availability requirement on the EC2 instance, Amazon RDS for MariaDB DB instance, and ElastiCache for Redis cluster

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: ADF

ADF, as mentioned in the other comments

upvoted 1 times

DmitriKonnovNN 1 year, 5 months ago

"The infrastructure can automatically recover from failure with the least possible downtime",

to me this sounds rather resilient than highly-availible, since it focuses on MITR but not explicitly on up-time.

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: ADF

obvious

upvoted 1 times

ansgohar 1 year, 9 months ago

Selected Answer: ADF

A, D, F

upvoted 1 times

NikkyDicky 2 years ago

Selected Answer: ADF

it's of course ADF

upvoted 1 times

Parimal1983 2 years ago

Selected Answer: ADF

For high availability, need to spin up instances in another zone with auto scaling and multi AZ options

upvoted 1 times

rtguru 2 years, 1 month ago

ADF will meet the described provisions

upvoted 1 times

RunkieMax 2 years, 1 month ago

Selected Answer: ADF

Fit the best the question

upvoted 1 times

Question #10 *Topic 1*

A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones.
After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs.
While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Choose two.)

A. Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3.

B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target.FailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.

C. Modify the existing Amazon Route 53 records by adding health checks. Configure a fallback target if the health check fails. Modify DNS records to point to a publicly accessible webpage.

D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.

E. Add a custom error response by configuring a CloudFront custom error page. Modify DNS records to point to a publicly accessible web page.

Suggested Answer: *CE*

Community vote distribution

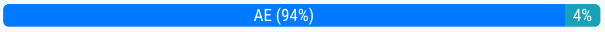

Raj40 Highly Voted 2 years, 6 months ago

A & E
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html#custom-error-pages-procedure
upvoted 35 times

atirado Highly Voted 9 months, 1 week ago

Selected Answer: AE

Option A - This option helps: Allows exposing custom error pages from a highly-available location

Option B - This option requires a lot of set up

Option C - This option might not work because modifying DNS will redirect all traffic publicly accessible webpage

Option D - This option requires a lot of set up

Option E - This option helps: Shows a custom error page when the error occurs
upvoted 9 times

5e8c031 Most Recent 1 week, 5 days ago

Selected Answer: AE

Any answer consisting of modifying the ALB forwarding rules because it will modify the behaviour of the system for all users, even those who did not get a HTTP 502. Any answer involving Route 53 will not work because it supposes that the DNS cache on clients expires before it takes effect, and it will effect all users, even those who did not get a HTTP 502.
Using a CloudFront custom error page is the only workable option that will serve the custome error page only in the event of a HTTP 502
upvoted 1 times

ausl 1 month, 3 weeks ago

Selected Answer: AE

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html#custom-error-pages-procedure
upvoted 1 times

teeee123 4 months, 1 week ago

Selected Answer: AE

Maybe A and E

upvoted 1 times

hoef03 8 months, 3 weeks ago

Selected Answer: AC

AC, beacuse for E to work, another s3 origin needs to be created, (OAI when not static website) , and a behaviour for the path routing to the error html.. Also unclear what the "DNS record motification" is made for.

upvoted 1 times

Parimal1983 9 months, 1 week ago

Selected Answer: AE

Custom error pages need to setup in different location then source (where web pages is hosted), configure CloudFront to use those custom error pages

upvoted 2 times

Sarutobi 9 months, 1 week ago

Selected Answer: AE

We need a combination, so A provides the error page; should we go with DNS health-check (C+A) or CloudFront (E+A)? In my case, I try to stick to a single service to do failover, and DNS is a great option, but it looks like, in this question, CloudFront is already present with the least-operational overhead.

upvoted 5 times

dev112233xx 9 months, 1 week ago

Selected Answer: AE

A&E are the correct answers imo

upvoted 1 times

amministrazione 10 months ago

A. Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3.

E. Add a custom error response by configuring a CloudFront custom error page. Modify DNS records to point to a publicly accessible web page.

upvoted 1 times

agatim 12 months ago

Selected Answer: AC

Option A - Allow us to expose a error page with low effort.

Option B - Requires a lot of set up

Option C - Allow us to redirect all the traffic to our error page exposed by S3 in case of errors.

Option D - requires a lot of set up

Option E - Custom Error Pages in CloudFront refers to the same Origin (in our case the Load Balancer) so it does not work with all the other answers.

So correct answer are A and C

upvoted 2 times

roger8t8 1 year ago

A & E

https://aws.amazon.com/blogs/aws/custom-error-pages-and-responses-for-amazon-cloudfront/

upvoted 1 times

azhar3128 1 year ago

I think it is wordplay. Option A says to upload "error pages", which will be an overhead for creating a page for each error and unnecessary. that's where C & E are correct

upvoted 4 times

shmoeee 5 months, 3 weeks ago

my exact thoughts. Why not use already available pages

upvoted 1 times

iulian0585 1 year, 1 month ago

Selected Answer: AE

A and E according to AWS documentation:

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html#custom-error-pages-procedur

upvoted 1 times

kz407 1 year, 3 months ago

Selected Answer: AE

The only problem with E is that it say "Modify DNS records to point to a publicly accessible web page" at the end. It doesn't make sense to begin with. And configuring custom error responses in CF has nothing to do with DNS anyway.

upvoted 6 times

MoT0ne 1 year, 3 months ago

I think why is not A is because of this sentence - "The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs." - so let's not think it as requiring a maintenance page

upvoted 1 times

abeb 1 year, 7 months ago

should be AE

upvoted 2 times

Question #11 *Topic 1*

A company has many AWS accounts and uses AWS Organizations to manage all of them. A solutions architect must implement a solution that the company can use to share a common network across multiple accounts.
The company's infrastructure team has a dedicated infrastructure account that has a VPC. The infrastructure team must use this account to manage the network. Individual accounts cannot have the ability to manage their own networks. However, individual accounts must be able to create AWS resources within subnets.
Which combination of actions should the solutions architect perform to meet these requirements? (Choose two.)

A. Create a transit gateway in the infrastructure account.

B. Enable resource sharing from the AWS Organizations management account.

C. Create VPCs in each AWS account within the organization in AWS Organizations. Configure the VPCs to share the same CIDR range and subnets as the VPC in the infrastructure account. Peer the VPCs in each individual account with the VPC in the infrastructure account.

D. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.

E. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each prefix list to associate with the resource share.

Suggested Answer: *AD*

Community vote distribution

BD (87%) 12%

masetromain Highly Voted 2 years, 6 months ago

Selected Answer: BD

I go with BD

upvoted 29 times

masetromain 2 years, 5 months ago

Step B is needed because it enables the organization to share resources across accounts.
Step D is needed because it allows the infrastructure account to share specific subnets with the other accounts in the organization, so that the other accounts can create resources within those subnets without having to manage their own networks.

upvoted 14 times

8693a49 11 months ago

Note that B says it enables sharing from the management account, but the infrastructure team must use the infrastructure account to manage the network", so there is nothing to share form the management account. Also, options D and E also enable resource sharing (you don't need to enable it from the management account, other accounts can enable resource sharing too).

VPCs can't talk to each other by default. You need to do something to 'glue' them together in a larger network.

upvoted 2 times

Kirkster 5 months, 2 weeks ago

In this case, there is actually only one VPC - the one in the infrastructure account. Users in other accounts deploy to subnets in that account, as those subnets are shared using resource sharing, as outlined in this blog:
https://aws.amazon.com/blogs/networking-and-content-delivery/vpc-sharing-a-new-approach-to-multiple-accounts-and-vpc-management/

upvoted 1 times

razguru Highly Voted 2 years, 6 months ago

A - Doesn't seem correct as the question didnt state multiple VPs, so transit gateway is not relevant.
I will go with B & D

upvoted 11 times

8693a49 11 months ago

There are multiple VPCs because each account must have at least one.

upvoted 3 times

Curious76 Most Recent 1 month, 2 weeks ago

Selected Answer: AD

A. Create a transit gateway in the infrastructure account.

AWS Transit Gateway allows multiple accounts to connect to a single shared network efficiently.

AWS Resource Access Manager (RAM) lets the infrastructure team share subnets with other accounts.
upvoted 1 times

unbornfroyo 2 months, 1 week ago

Selected Answer: AD

✓ A. Transit Gateway in Infrastructure Account
A transit gateway (TGW) allows centralized, scalable network connectivity between VPCs across accounts.

Hosting the TGW in the infrastructure account ensures centralized control of routing and traffic flow, while allowing other accounts to connect through attachments.

✓ D. RAM Share of Subnets
Use AWS Resource Access Manager (RAM) in the infrastructure account to share subnets with specific AWS accounts or OUs in the organization.

This allows other accounts to launch resources (like EC2 instances) into shared subnets without being able to modify the network itself.
upvoted 1 times

Kirkster 5 months, 2 weeks ago

Selected Answer: BD

The transit gateway is a red herring. Creating a resource share and then sharing out subnets is described in this AWS blog:
https://aws.amazon.com/blogs/networking-and-content-delivery/vpc-sharing-a-new-approach-to-multiple-accounts-and-vpc-management/

To those who say there are multiple VPCs (one in each account), read the question more carefully - it never says that. It merely says that users in the other accounts need to be able to deploy resources to the shared subnets.
upvoted 1 times

_KBM 6 months, 1 week ago

Selected Answer: AD

A. Create a transit gateway in the infrastructure account.

A transit gateway allows the infrastructure account to centralize the network and connect multiple VPCs across accounts. It serves as the backbone for communication between the accounts.
D. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.

Using AWS Resource Access Manager (RAM), you can share the subnets of the infrastructure account's VPC with other accounts in the organization, enabling individual accounts to create resources in those subnets while centralizing network management in the infrastructure account.
upvoted 1 times

wem 6 months, 3 weeks ago

Selected Answer: AD

The following steps will meet the requirements for sharing a common network managed from the infrastructure account across multiple AWS accounts, with the least operational complexity and in line with best practices:

A. Create a transit gateway in the infrastructure account.

A transit gateway allows centralized routing and connectivity between multiple VPCs across different AWS accounts. This approach enables the infrastructure account to control network management while allowing other accounts to use the shared network.
D. Create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.

Using AWS Resource Access Manager (RAM) to share subnets from the infrastructure account allows individual accounts to create resources within those subnets. This aligns with the requirement that individual accounts can create resources but not manage the network.
upvoted 1 times

Heman31in 6 months, 3 weeks ago

Selected Answer: AD

Without Step A (transit gateway), the solution would lack a central mechanism for connecting VPCs across accounts, which is essential for a shared network. D because: AWS Resource Access Manager (RAM) allows you to share VPC subnets across accounts. By creating a resource share for specific subnets and associating it with the appropriate organizational units (OUs), individual accounts can launch resources in the shared subnets while the infrastructure account retains network control.

upvoted 1 times

TariqKipkemei 7 months, 3 weeks ago

Selected Answer: BD

Enable resource sharing from the AWS Organizations management account then,
create a resource share in AWS Resource Access Manager in the infrastructure account. Select the specific AWS Organizations OU that will use the shared network. Select each subnet to associate with the resource share.

upvoted 1 times

Sin_Dan 8 months, 2 weeks ago

The correct answer is A and D.
B is a wrong option.

While AWS Organizations is required to manage multiple accounts, enabling resource sharing through AWS RAM is done in the infrastructure account (where the VPC resides), not in the AWS Organizations management account. Resource sharing is configured via RAM in the account that owns the resources, not through Organizations directly.

upvoted 2 times

SkyZeroZx 9 months, 1 week ago

Selected Answer: BD

The correct answers are D and B.

D will allow the infrastructure team to create a resource share in AWS Resource Access Manager in the infrastructure account. This will allow them to share the VPC with the other accounts in the organization.

B will enable resource sharing from the AWS Organizations management account. This is required to allow the resource share to be created.

C is not necessary, as the resource share will allow the other accounts to create resources in the shared VPC.

A is not necessary, as the resource share will allow the other accounts to connect to the shared VPC through the transit gateway.

E is not necessary, as the resource share will allow the other accounts to create resources in the shared VPC without the need for prefix lists.

upvoted 1 times

sreed77 9 months, 1 week ago

Selected Answer: BD

Option B allows the infrastructure team to manage the network in the infrastructure account. It also allows individual accounts to create AWS resources within subnets. This is done by creating a resource share in AWS Resource Access Manager (RAM) in the infrastructure account. The resource share is then associated with the specific AWS Organizations OU that will use the shared network. The subnets are then associated with the resource share.

Option D is also necessary because it allows the infrastructure team to control who has access to the shared network. This is done by assigning permissions to the resource share.

Here are the steps involved in implementing this solution:

Create a resource share in RAM in the infrastructure account.
Select the specific AWS Organizations OU that will use the shared network.
Select each subnet to associate with the resource share.
Assign permissions to the resource share.

upvoted 4 times

cnethers 9 months, 1 week ago

I would go BD
When you share a subnet using AWS Resource Access Manager (RAM) with another AWS account, the resources within that shared subnet can communicate with each other and with the resources in the account that owns the subnet. However, for outbound network connectivity to other VPCs, on-premises networks, or the internet, you need to set up additional networking components.

upvoted 1 times

cnethers 9 months, 1 week ago

2. Inter-VPC Communication:

o If the resources in the shared subnet need to communicate with resources in another VPC (either within the same AWS account or in a different AWS account), you can use VPC Peering or a Transit Gateway.

o VPC Peering: Establish a peering connection between the VPCs and update the route tables accordingly.

o Transit Gateway: Create a Transit Gateway, attach both VPCs to the Transit Gateway, and configure the necessary route tables and Transit Gateway route tables.

upvoted 1 times

cnethers 1 year ago

Here's a breakdown of different scenarios and the required setup:

1. Internet Access:

o If you need resources in the shared subnet to access the internet, ensure that the subnet is a public subnet with an associated Internet Gateway (IGW) and appropriate route table entries.

o The account that owns the VPC will typically manage the IGW and the route tables.

upvoted 1 times

cnethers 1 year ago

3. On-Premises Connectivity:

o If the resources in the shared subnet need to communicate with an on-premises network, you can use AWS Direct Connect or a Site-to-Site VPN.

o These connections can be routed through a Transit Gateway for more scalable and manageable network architecture.

upvoted 1 times

shaaam80 9 months, 1 week ago

Selected Answer: BD

Answer - B & D.

A is wrong. No TGW needed as customer has just 1 VPC.

E is wrong - can't share resources via RAM using prefix lists.

C is wrong - talks about creating VPCs with same CIDR ranges and VPC peering (not possible with overlapping CIDRs and not needed for this solution as there is just 1 VPC).

upvoted 3 times

Sin_Dan 8 months, 2 weeks ago

How do you think the Accounts got subnets without VPCs?

upvoted 1 times

severlight 9 months, 1 week ago

Selected Answer: BD

I don't see the way you can share a prefix list.

upvoted 2 times

mattfaz 8 months, 2 weeks ago

https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html

upvoted 1 times

8693a49 11 months ago

You don't share a prefix list, you associate it with the shared resource (which here is a TGW). The way you do it is you add the prefixes to the route tables inside the account's VPCs. The prefixes will point towards the TGW. This makes the network traffic destined to other account go through the TGW into these accounts based on the TGW routing table. The TGW routing table can only be controlled from the infrastructure account.

upvoted 2 times

AlbertS82 9 months, 1 week ago

Selected Answer: BD

B&D is the only correct answer

upvoted 3 times

atirado 9 months, 1 week ago

Selected Answer: BD

Option A - Does not assist with allowing OUs to create resources in the subnets

Option B - Allows sharing resources across the entire organization

Option C - This option does not work as a way to share subnets because it creates multiple VPCs and subnets in the accounts rather than allowing managing resources in shared subnets

Option D - Directly shares the subnets

Option E - Does not assist because it only shares pre-built CIDR blocks rather than subnets
upvoted 4 times

8693a49 11 months ago
Subnets cannot be shared
upvoted 1 times

Question #12 *Topic 1*

A company wants to use a third-party software-as-a-service (SaaS) application. The third-party SaaS application is consumed through several API calls. The third-party SaaS application also runs on AWS inside a VPC.
The company will consume the third-party SaaS application from inside a VPC. The company has internal security policies that mandate the use of private connectivity that does not traverse the internet. No resources that run in the company VPC are allowed to be accessed from outside the company's VPC. All permissions must conform to the principles of least privilege.
Which solution meets these requirements?

A. Create an AWS PrivateLink interface VPC endpoint. Connect this endpoint to the endpoint service that the third-party SaaS application provides. Create a security group to limit the access to the endpoint. Associate the security group with the endpoint.

B. Create an AWS Site-to-Site VPN connection between the third-party SaaS application and the company VPC. Configure network ACLs to limit access across the VPN tunnels.

C. Create a VPC peering connection between the third-party SaaS application and the company VPUpdate route tables by adding the needed routes for the peering connection.

D. Create an AWS PrivateLink endpoint service. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint service. Grant permissions for the endpoint service to the specific account of the third-party SaaS provider.

Suggested Answer: *A*

Community vote distribution

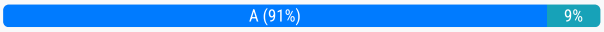

Raj40 Highly Voted 2 years, 6 months ago

Selected Answer: A

https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html

upvoted 19 times

masetromain Highly Voted 2 years, 6 months ago

Selected Answer: A

I go with A

upvoted 8 times

masetromain 2 years, 5 months ago

A. Create an AWS PrivateLink interface VPC endpoint. Connect this endpoint to the endpoint service that the third-party SaaS application provides. Create a security group to limit the access to the endpoint. Associate the security group with the endpoint.
This solution uses AWS PrivateLink, which creates a secure and private connection between the company's VPC and the third-party SaaS application VPC, without the traffic traversing the internet. The use of a security group and limiting access to the endpoint service conforms to the principle of least privilege.

upvoted 11 times

2aldous Most Recent 9 months, 1 week ago

Selected Answer: A

Access Saas products throgh AWS Private Link is the answer.

upvoted 1 times

SkyZeroZx 9 months, 1 week ago

Selected Answer: A

Create an AWS PrivateLink interface VPC endpoint.

upvoted 1 times

NikkyDicky 9 months, 1 week ago

Selected Answer: A

it s a

upvoted 1 times

cattle_rei 9 months, 1 week ago

Selected Answer: A

It's A because in this scenario we are consuming a service , not providing one, so that eliminates E .

upvoted 1 times

shaaam80 9 months, 1 week ago

Selected Answer: A

Answer - A.

VPC Interface end point to access any service privately without traversing the internet.

AWS Private Link VPC endpoint to access the SaaS application.

upvoted 1 times

atirado 9 months, 1 week ago

Selected Answer: A

Option A - The interface VPC Endpoint will provide local access to the SaaS service from within the company's VPC. Moreover, traffic to and access from the SaaS VPC will traverse the AWS network rather than the internet. This is considered private traffic.

Option B - This option might not work: Nothing is said about whether the CIDR blocks in each VPC overlap. Moreover, nothing is said about whether bandwidth limitations on Site-Site VPN could be an issue.

Option C - This option might not work: Nothing is said about whether the CIDR blocks in each VPC overlap.

Option D - This option will not work: A PrivateLink Endpoint service is used for facilitating access to AWS services.

upvoted 3 times

gofavad926 9 months, 1 week ago

Selected Answer: A

A, the service provider creates an endpoint service and grants their customers access to the endpoint service. As the service consumer, you create an interface VPC endpoint, which establishes connections between one or more subnets in your VPC and the endpoint service.

upvoted 1 times

amministrazione 10 months ago

A. Create an AWS PrivateLink interface VPC endpoint. Connect this endpoint to the endpoint service that the third-party SaaS application provides. Create a security group to limit the access to the endpoint. Associate the security group with the endpoint.

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: A

obvious

upvoted 1 times

senthilsekaran 1 year, 8 months ago

Correct Answer : A

upvoted 1 times

task_7 1 year, 9 months ago

Selected Answer: D

A VS D

A. Create an AWS PrivateLink interface VPC endpoint. Connect this endpoint to the endpoint service that the third-party SaaS application provides.

D. Create an AWS PrivateLink endpoint service. Ask the third-party SaaS provider to create an interface VPC endpoint for this endpoint service

D is right SaaS provider has create interface VPC endpoint for this endpoint service

upvoted 4 times

Jassybanga 1 year, 4 months ago

exactly , we need to access the resource from SAAS Provider and not vice versa , Hence in this case the VPC Gateway endpoint should be provided from SAAS Provider for the privatelink endpoint we provide it to them - we use this for Snowflake Saas :)

upvoted 1 times

whenthan 1 year, 10 months ago

Selected Answer: A

https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html

https://aws.amazon.com/blogs/apn/enabling-new-saas-strategies-with-aws-privatelink/

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: A

Create an AWS PrivateLink interface VPC endpoint.

upvoted 1 times

kiran15789 2 years, 3 months ago

Selected Answer: A

https://docs.aws.amazon.com/vpc/latest/privatelink/privatelink-access-saas.html

upvoted 1 times

ptpho 2 years, 6 months ago

It's A .clearly

upvoted 4 times

Question #13 *Topic 1*

A company needs to implement a patching process for its servers. The on-premises servers and Amazon EC2 instances use a variety of tools to perform patching. Management requires a single report showing the patch status of all the servers and instances.
Which set of actions should a solutions architect take to meet these requirements?

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instances. Use Systems Manager to generate patch compliance reports.

B. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances. Use Amazon QuickSight integration with OpsWorks to generate patch compliance reports.

C. Use an Amazon EventBridge rule to apply patches by scheduling an AWS Systems Manager patch remediation job. Use Amazon Inspector to generate patch compliance reports.

D. Use AWS OpsWorks to manage patches on the on-premises servers and EC2 instances. Use AWS X-Ray to post the patch status to AWS Systems Manager OpsCenter to generate patch compliance reports.

Suggested Answer: *A*

Community vote distribution

A (100%)

masetromain Highly Voted 2 years, 6 months ago

Selected Answer: A

A is good
https://docs.aws.amazon.com/prescriptive-guidance/latest/patch-management-hybrid-cloud/design-on-premises.html
upvoted 14 times

masetromain 2 years, 5 months ago

A is correct. AWS Systems Manager can manage patches on both on-premises servers and EC2 instances and can generate patch compliance reports. AWS OpsWorks and Amazon Inspector are not specifically designed for patch management and therefore would not be the best choice for this use case. Using Amazon EventBridge rule and AWS X-Ray to generate patch compliance reports is not a practical solution as they are not designed for patch management reporting.
upvoted 13 times

TariqKipkemei Most Recent 7 months, 3 weeks ago

Selected Answer: A

Use AWS Systems Manager update, manage, and configure Amazon EC2 instances, edge devices, on-premises servers, and virtual machines (VMs).
upvoted 1 times

atirado 9 months, 1 week ago

Selected Answer: A

Option A - Systems Manager patches and generates patch compliance reports.

Option B - This option does not apply because Chef or Puppet are not mentioned in the question. Moreover, either one does not directly perform patch management.

Option C - Inspector would generate a report for on-premise resources

Option D - This option does not apply because Chef or Puppet are not mentioned in the question. Moreover, X-Ray does apply.
upvoted 1 times

MoT0ne 9 months, 1 week ago

Selected Answer: A

AWS OpsWorks is a configuration management service that provides a way to automate the deployment, configuration, and management of applications on EC2 instances. It is designed to help you manage the entire lifecycle of your applications.
upvoted 1 times

severlight 9 months, 1 week ago

Selected Answer: A

obvious

upvoted 1 times

whenthan 9 months, 1 week ago

Selected Answer: A

A is correct

upvoted 1 times

amministrazione 10 months ago

A. Use AWS Systems Manager to manage patches on the on-premises servers and EC2 instances. Use Systems Manager to generate patch compliance reports.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A is the correct answer

upvoted 1 times

stevegod0 1 year, 11 months ago

A is correct:

https://www.amazonaws.cn/en/systems-manager/

upvoted 1 times

cattle_rei 1 year, 11 months ago

Selected Answer: A

Other options are distractors. Opswork would be right only if customer wanted to make use of existing script or know-how in chef or puppet.

upvoted 1 times

NikkyDicky 2 years ago

Selected Answer: A

yep - A

upvoted 1 times

EricZhang 2 years, 1 month ago

A is the best but Systems Manager cannot generate the patch compliance reports.

https://docs.aws.amazon.com/prescriptive-guidance/latest/patch-management-hybrid-cloud/design-on-premises.html

- A resource data sync in Systems Manager Inventory gathers the patching details and publishes them to an S3 bucket.

- Patch compliance reporting and dashboards are built in Amazon QuickSight from the S3 bucket information.

upvoted 1 times

gameoflove 2 years, 1 month ago

Selected Answer: A

A is the right answer for this question as per information shared by them

upvoted 2 times

2aldous 2 years, 2 months ago

Selected Answer: A

Easy question :)

A is the answer.

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: A

Use AWS Systems Manager to manage patches

upvoted 1 times

kiran15789 2 years, 3 months ago

Selected Answer: A

https://docs.aws.amazon.com/prescriptive-guidance/latest/patch-management-hybrid-cloud/design-on-premises.html

upvoted 1 times

gameoflove 2 years, 4 months ago

Selected Answer: A

AWS System Manager support On-premise and EC2 instance patching

upvoted 2 times

Question #14 *Topic 1*

A company is running an application on several Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The load on the application varies throughout the day, and EC2 instances are scaled in and out on a regular basis. Log files from the EC2 instances are copied to a central Amazon S3 bucket every 15 minutes. The security team discovers that log files are missing from some of the terminated EC2 instances.
Which set of actions will ensure that log files are copied to the central S3 bucket from the terminated EC2 instances?

A. Create a script to copy log files to Amazon S3, and store the script in a file on the EC2 instance. Create an Auto Scaling lifecycle hook and an Amazon EventBridge rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to send ABANDON to the Auto Scaling group to prevent termination, run the script to copy the log files, and terminate the instance using the AWS SDK.

B. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.

C. Change the log delivery rate to every 5 minutes. Create a script to copy log files to Amazon S3, and add the script to EC2 instance user data. Create an Amazon EventBridge rule to detect EC2 instance termination. Invoke an AWS Lambda function from the EventBridge rule that uses the AWS CLI to run the user-data script to copy the log files and terminate the instance.

D. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic. From the SNS notification, call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send ABANDON to the Auto Scaling group to terminate the instance.

Suggested Answer: *B*

Community vote distribution

masetromain Highly Voted 9 months, 1 week ago

Selected Answer: B

B. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance. This approach will use the Auto Scaling lifecycle hook to execute the script that copies log files to S3, before the instance is terminated, ensuring that all log files are copied from the terminated instances.

upvoted 12 times

rtgfdv3 Highly Voted 2 years, 6 months ago

Selected Answer: B

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

upvoted 8 times

pk0619 Most Recent 6 months, 2 weeks ago

Selected Answer: B

This is most accurate with redundancy as EventBridge can directly invoke SSM Document and you don't need Lambda function.

upvoted 1 times

Shanmahi 7 months ago

Selected Answer: B

B using systems manager

upvoted 1 times

atirado 9 months, 1 week ago

Selected Answer: B

Option A - This option might not work: Preventing ASG termination could create further trouble and there is no guarantee the script will run if the instance happens to be unhealthy

Option B - This option could work: Running the script from the SSM API guarantees the script will run, using EventBridge to capture the ASG

termination event provides a perfect place to hook in the call to SSM which will also pause the termination until the script runs. Then CONTINUE allows the ASG termination to continue.

Option C - This option does not work because it does not solve the problem: Terminating instances within the 15 minute window causes log files to be lost.

Option D - This option might not work: It does not rely on EventBridge to detect the ASG termination event. It also could create further trouble because no other actions will be performed due to sending ABANDON though nothing is said about other actions in the question

upvoted 6 times

F_Eldin 9 months, 1 week ago

Selected Answer: B

A- Wrong because prevent termination is not needed.

C- Wrong because 5-minute frequency creates an overhead or delay . Using user data for the script adds complexity

D- Wrong because SNS

upvoted 3 times

gameoflove 9 months, 1 week ago

Selected Answer: B

B is the right answer due to Auto Scaling lifecycle hook and an Amazon EventBridge rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send

upvoted 1 times

cattle_rei 9 months, 1 week ago

Selected Answer: B

I think this is B. It could be A as well, but B is better solution because the document with SM can be re-utilized with other instances. Also A would require using a custom image with the script or user data to create the script, so more points of failure.

upvoted 1 times

ansgohar 9 months, 1 week ago

Selected Answer: B

B. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.

upvoted 1 times

amministrazione 10 months ago

B. Create an AWS Systems Manager document with a script to copy log files to Amazon S3. Create an Auto Scaling lifecycle hook and an Amazon EventBridge rule to detect lifecycle events from the Auto Scaling group. Invoke an AWS Lambda function on the autoscaling:EC2_INSTANCE_TERMINATING transition to call the AWS Systems Manager API SendCommand operation to run the document to copy the log files and send CONTINUE to the Auto Scaling group to terminate the instance.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B is the correct answer

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: B

both abandon and continue will lead to instance termination, the difference is abandon will prevent from running other lifycycle hooks

upvoted 2 times

cattle_rei 1 year, 10 months ago

I think this is B. It could be A as well, but B is better solution because the document with SM can be re-utilized with other instances. Also A would require using a custom image with the script or user data to create the script, so more points of failure.

upvoted 1 times

softarts 1 year, 10 months ago

Selected Answer: B

d is wrong, shouldn't be "ABANDON"

upvoted 2 times

NikkyDicky 2 years ago

Selected Answer: B

it's a B

upvoted 1 times

2aldous 2 years, 2 months ago

Selected Answer: B

B.

Smart solution :)

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: B

Systems manager + eventbridge

upvoted 4 times

Question #15 *Topic 1*

A company is using multiple AWS accounts. The DNS records are stored in a private hosted zone for Amazon Route 53 in Account A. The company's applications and databases are running in Account B.
A solutions architect will deploy a two-tier application in a new VPC. To simplify the configuration, the db.example.com CNAME record set for the Amazon RDS endpoint was created in a private hosted zone for Amazon Route 53.
During deployment, the application failed to start. Troubleshooting revealed that db.example.com is not resolvable on the Amazon EC2 instance.
The solutions architect confirmed that the record set was created correctly in Route 53.
Which combination of steps should the solutions architect take to resolve this issue? (Choose two.)

A. Deploy the database on a separate EC2 instance in the new VPC. Create a record set for the instance's private IP in the private hosted zone.

B. Use SSH to connect to the application tier EC2 instance. Add an RDS endpoint IP address to the /etc/resolv.conf file.

C. Create an authorization to associate the private hosted zone in Account A with the new VPC in Account B.

D. Create a private hosted zone for the example com domain in Account B. Configure Route 53 replication between AWS accounts.

E. Associate a new VPC in Account B with a hosted zone in Account A. Delete the association authorization in Account A.

Suggested Answer: *BC*

Community vote distribution

CE (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: CE

C and E are correct.

C. Create an authorization to associate the private hosted zone in Account A with the new VPC in Account B.
This step is necessary because the VPC in Account B needs to be associated with the private hosted zone in Account A to be able to resolve the DNS records.

E. Associate a new VPC in Account B with a hosted zone in Account A. Delete the association authorization in Account A.
This step is necessary because the association authorization needs to be removed in Account A after the association is done in Account B.
upvoted 35 times

kiran15789 Highly Voted 2 years, 3 months ago

Selected Answer: CE

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/hosted-zone-private-associate-vpcs-different-accounts.html
upvoted 10 times

TariqKipkemei Most Recent 7 months, 3 weeks ago

Selected Answer: CE

Associate the new VPC in Account B with the hosted zone in Account A, delete the association authorization in Account A.
Then create an authorization to associate the private hosted zone in Account A with the new VPC in Account B.
upvoted 1 times

masetromain 9 months, 1 week ago

Selected Answer: CE

With comments and links the answer is C and E. (Ty robertohyène and JosuéXu)

C = 6. Run the following command to create the association between Account A's private hosted zone and Account B's VPC. Use the hosted zone's ID from step 3. B account.
E = 7. It is recommended to remove the association permission after the association is created. This will prevent you from recreating the same association later.

https://aws.amazon.com/premiumsupport/knowledge-center/route53-private-hosted-zone/
upvoted 4 times

masetromain 2 years, 6 months ago

https://www.examtopics.com/discussions/amazon/view/36113-exam-aws-certified-solutions-architect-professional-topic-1/

upvoted 1 times

CloudFloater 9 months, 1 week ago

Selected Answer: CE

C and E.

In order to resolve the issue, the solutions architect should create an authorization to associate the private hosted zone in Account A with the new VPC in Account B (Option C). This will allow the new VPC in Account B to access the DNS records stored in the private hosted zone in Account A.

In addition, the solutions architect should associate the new VPC in Account B with the hosted zone in Account A (Option E) and delete the association authorization in Account A. This will ensure that the new VPC in Account B is properly configured to use the private hosted zone in Account A and resolve the db.example.com CNAME record set correctly.

upvoted 5 times

whenthan 9 months, 1 week ago

Selected Answer: CE

https://repost.aws/knowledge-center/route53-private-hosted-zone

Create an authorization to associate the private hosted zone and as a best practice , it is recommended to delete the association authorization in account A-This step prevents you from recreating the same association later. To delete the authorization, reconnect to the EC2 instance in Account A

upvoted 2 times

liuliangzhou 9 months, 2 weeks ago

Selected Answer: CE

A account's DNS Zone authorization is associated with B's VPC, and after B's VPC is associated with A's Priviate Zone, A's authorization permission is deleted for security reasons.

upvoted 1 times

amministrazione 10 months ago

C. Create an authorization to associate the private hosted zone in Account A with the new VPC in Account B.

E. Associate a new VPC in Account B with a hosted zone in Account A. Delete the association authorization in Account A.

upvoted 1 times

7f6aef3 1 year, 2 months ago

Selected Answer: CE

https://repost.aws/knowledge-center/route53-private-hosted-zone

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: CE

Correct answers

upvoted 1 times

8608f25 9 months, 1 week ago

Explanation:

* Option C is correct because, in a multi-account AWS setup, to use a Route 53 private hosted zone from one account (Account A) in another account's VPC (Account B), you first need to create an authorization. This authorization is necessary for allowing the private hosted zone in one account to be associated with a VPC in another account. This step enables the resolution of DNS records stored in the private hosted zone across accounts.

* Option E is correct as it follows up on the authorization created in Option C. Once the authorization is in place, you can then associate the new VPC in Account B with the private hosted zone in Account A. This association is what actually allows the EC2 instances within the VPC in Account B to resolve DNS queries using the private hosted zone in Account A, ensuring that db.example.com can be resolved as intended.

upvoted 4 times

8608f25 1 year, 4 months ago

Why the others are incorrect:

* Option A is not a direct solution to the problem of DNS resolution across AWS accounts. Deploying the database on an EC2 instance does not address the issue of DNS resolution for the RDS endpoint across accounts.

* Option B is not a scalable or AWS-recommended solution. Manually adding RDS endpoint IP addresses to the /etc/resolv.conf file on an EC2 instance is not practical for environments that require automation and could lead to issues if the RDS endpoint changes.

* Option D involves creating a separate private hosted zone in Account B and configuring Route 53 replication between AWS accounts. This option is unnecessary and more complex than required. The direct association of VPCs across accounts to a single hosted zone is a simpler and more effective solution.

Therefore, Options C and E are the steps that directly address the issue with the least complexity and enable the intended DNS resolution across AWS accounts.

upvoted 3 times

atirado 1 year, 6 months ago

Selected Answer: CE

Option A - This option does not work - It does not provide for solving address name resolution in the new VPC

Option B - This option works but it breaks the company's architecture where all DNS names are stored in the private zone in Account A

Option C - This option contributes to the solution.

Option D - Breaks the company's architecture

Option E - This option contributes to the solution

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: CE

obvious

upvoted 1 times

SfQ 1 year, 8 months ago

Selected Answer: CE

C and E are correct.

B is not a best solution. It's a manual setup and it may lose the configuration if we are using ASG and launching new instance.

upvoted 1 times

Chainshark 1 year, 8 months ago

Why is B marked as correct?

upvoted 2 times

SfQ 1 year, 8 months ago

B is not a best solution. It's a manual setup and it may lose the configuration if we are using ASG and launching new instance.

upvoted 2 times

NikkyDicky 2 years ago

Selected Answer: CE

it's CE

upvoted 1 times

Jonalb 2 years ago

Selected Answer: CE

ccccccccccccceeeeeeeeeeeee

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: CE

C & E as Issue is associated with authorization

upvoted 1 times

Question #16 *Topic 1*

A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.
Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

A. Reconfigure Amazon EFS to enable maximum I/O.

B. Update the blog site to use instance store volumes for storage. Copy the site contents to the volumes at launch and to Amazon S3 at shutdown.

C. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.

D. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

Suggested Answer: *C*

Community vote distribution

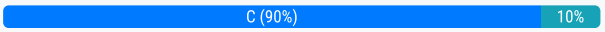

masetromain Highly Voted 9 months, 1 week ago

Selected Answer: C

C. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.

Amazon CloudFront is a content delivery network (CDN) that can be used to deliver content to users with low latency and high data transfer speeds. By configuring a CloudFront distribution for the blog site and pointing it at an S3 bucket, the videos can be cached at edge locations closer to users, reducing buffering and timeout issues. Additionally, S3 is designed for scalable storage and can handle high levels of user traffic. Migrating the videos from EFS to S3, would also improve the performance and scalability of the website.

upvoted 28 times

spencer_sharp Highly Voted 2 years, 6 months ago

Selected Answer: C

No brainer

upvoted 9 times

TariqKipkemei Most Recent 7 months, 3 weeks ago

Selected Answer: C

video ,unstructured content = Amazon S3
resolve buffering and timeout issues = Amazon CloudFront distribution

upvoted 1 times

ninomfr64 9 months, 1 week ago

Selected Answer: C

Not A as Max I/O increase IOPS but negatively impact latency, ultimately you will have little to no performance improvement. Also you cannot enable Max IO on an existing filesystem.
Not B as this is not a cheap option (instance store generally cost more than EBS backed), also without a CDN there will be little performance improvement
Not D as this provides performance improvements, but this provide comparable performance to option C at higher costs as in D videos are stored on EFS that cost more than S3 and all traffic goes trough CDN rather than only videos that actually needs eddge caching

Thus C provide performance improvements (thanks for CloudFront) with cost-effective approach (S3 is cheap)

upvoted 2 times

ninomfr64 1 year, 6 months ago

Also this follows AWS best practices to separate static content from dynamic content allowing for better scalability

upvoted 1 times

atirado 9 months, 1 week ago

Selected Answer: C

Option A - This option might not work and is not cheap: It will increase costs and has limited scalability. EFS is an expensive storage solution for videos

Option B - This option might not work: Nothing is mentioned about whether the application is stateful or stateless and whether the ALB has client stickiness so using instance store could provide an inconsistent user experience. S3 is a cheap storage option

Option C - This option will work and is cheap: A CloudFront distribution and S3 will provide the most scalability and availability possible from AWS; and both are very cheap options for distribution and storage of content

Option D - This option might work but is not cheap: Moving all content to CloudFront ensures it will be served from the edge cache for the duration of the cache mitigating issues during high usage. However, nothing is said in the question about usage patterns, i.e performance issue will happen again for older content. Moreover, EFS is an expensive storage solution for video files compared to S3.

upvoted 1 times

amministrazione 10 months ago

C. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.

upvoted 1 times

Bereket 1 year ago

Selected Answer: D

The most cost-efficient and scalable deployment that will resolve the issues for users, given the requirements and the described scenario, is:

D. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.

upvoted 2 times

Christophe_ 1 year, 4 months ago

Selected Answer: D

Option C - Does not support new content added later by users, does not accelerate site content

Option D - Accelerate site and videos, allow content added

upvoted 2 times

e4bc18e 1 year, 3 months ago

Cloudfront caches data to serve more rapidly at the edge and not have to serve content from the backend, that is acceleration. Also you can now write to S3 for new data. Sorry your choice is not correct.

upvoted 2 times

geekos 1 year, 7 months ago

Selected Answer: C

C is good

upvoted 1 times

abeb 1 year, 7 months ago

C is good

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: C

obvious

upvoted 1 times

cattle_rei 1 year, 10 months ago

Selected Answer: C

No doubt it's C. To me the keyword there is scalable. S3 will be able to handle any amount of content users can generate. EFS is not the right solution for object storage, s3 is. EFS is a solution for a sharable network filesystem, that can be mounted and used by many operation systems.

upvoted 1 times

Magoose 1 year, 11 months ago

Selected Answer: D

C and D are both viable. But D would be less overhead as you would most likely need to reconfigure the web application more to get it working with S3. Option D with Elastic Beanstalk provides a higher level of abstraction and automates many aspects of the application management, which can reduce operational overhead and simplify the re-architecting process

upvoted 2 times

totopopo 1 year, 11 months ago

D is not cost effective, which was the demand for the question.
If it was about less changes, I would go with it.
Here, right answer is C.
upvoted 1 times

NikkyDicky 2 years ago

C more cost efficient
upvoted 1 times

karim_arous 2 years ago

Selected Answer: C

C without a doubt
upvoted 1 times

gameoflove 2 years, 1 month ago

Selected Answer: C

C is only option which meet their requirement
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

Configure an Amazon CloudFront distribution.
upvoted 2 times

Question #17 *Topic 1*

A company with global offices has a single 1 Gbps AWS Direct Connect connection to a single AWS Region. The company's on-premises network uses the connection to communicate with the company's resources in the AWS Cloud. The connection has a single private virtual interface that connects to a single VPC.
A solutions architect must implement a solution that adds a redundant Direct Connect connection in the same Region. The solution also must provide connectivity to other Regions through the same pair of Direct Connect connections as the company expands into other Regions.
Which solution meets these requirements?

A. Provision a Direct Connect gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the Direct Connect gateway. Connect the Direct Connect gateway to the single VPC.

B. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new private virtual interface on the new connection, and connect the new private virtual interface to the single VPC.

C. Keep the existing private virtual interface. Create the second Direct Connect connection. Create a new public virtual interface on the new connection, and connect the new public virtual interface to the single VPC.

D. Provision a transit gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the transit gateway. Associate the transit gateway with the single VPC.

Suggested Answer: *A*

Community vote distribution

A (100%)

masetromain Highly Voted 9 months, 1 week ago

Selected Answer: A

A. Provision a Direct Connect gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the Direct Connect gateway. Connect the Direct Connect gateway to the single VPC.

This solution provides a redundant Direct Connect connection in the same Region by creating a new private virtual interface on each connection, and connecting both private virtual interfaces to a Direct Connect gateway. The Direct Connect gateway is then connected to the single VPC. This solution also allows the company to expand into other Regions while providing connectivity through the same pair of Direct Connect connections.
The Direct Connect Gateway allows you to connect multiple VPCs and on-premises networks in different accounts and different regions to a single Direct Connect connection.
It also provides automatic failover and routing capabilities.
upvoted 24 times

masetromain 2 years, 5 months ago

Option D is not the best solution because it uses a Transit Gateway, which is used to connect multiple VPCs and on-premises networks in different accounts and different regions, but it is not necessary in this scenario. The company only wants to add a redundant Direct Connect connection in the same Region and connect it to the same VPC. Additionally, using a Transit Gateway in this scenario would add more complexity and might not be necessary.
Also, Transit Gateway does not provide automatic failover and routing capabilities, which is required in this scenario.
The Direct Connect Gateway is a better choice in this scenario as it provides the necessary functionality of automatic failover and routing capabilities, and it is more suitable for connecting multiple Direct Connect connections to a single VPC.
upvoted 16 times

Sarutobi 2 years, 4 months ago

All options here are problematic. The DX-GW is a control plane-only device; in other words, no actual traffic goes over it; it is just a Route-Reflector it only carries the routing table. TGW is not a region construct, so by itself, it cannot provide regional redundancy. In any case, all things considered, maybe A is the closest but it should mention VGW.
upvoted 2 times

Sarutobi 2 years, 4 months ago

I meant to say, "TGW is a region construct".

upvoted 1 times

anita_student 2 years, 4 months ago

Option D is not possible at all. You connect to TGW using transit VIF, not private VIF

upvoted 9 times

AMohanty 1 year, 9 months ago

Transit GW - connects both over Private VIF and Transit VIF

upvoted 1 times

kz407 Highly Voted 1 year, 3 months ago

What I don't understand is why do you need to delete the existing private VIF? Can't that be reassigned?

upvoted 6 times

TariqKipkemei Most Recent 7 months, 3 weeks ago

Selected Answer: A

'must provide connectivity to other Regions through the same pair of Direct Connect connections'= Direct Connect gateway

upvoted 1 times

atirado 9 months, 1 week ago

Selected Answer: A

Option A - This option might work however it is missing a step: Connecting the Direct Connect Gateway to a Virtual Private Gateway in the single VPC (and any VPC in a new region)

Option B - This option will not work: It does not allow to grow into new regions and it does not create a redundant link

Option C - This option will not work: Using a Public Virtual interface does not connect VPC resources to on-premise

Option D - This option might work however it missing multiple steps: Each VPC will require its own Transit Gateway. Each Transit Gateway will connect through an association with Direct Connect gateway. Each Direct Connect connection will connect to the Direct Connect Gateway using a Transit VIF

upvoted 2 times

ninomfr64 9 months, 1 week ago

Selected Answer: A

I have to admit that initially I picked a wrong answer, here is my findings after some docs browsing:
Not B as this will provide Direct Connect (DX) redundancy but does not provide connectivity to other Regions
Not C as this will not even provide DX redundancy for the VPC because the public VIF on the new connection does not provide access to the VPC
Not D as Transit Gateway (TGW) is a regional resources and does not allows to provide connectivity to other Regions (you can peer with a TGW in another Region). Also you need to have a Transit virtual interface to connect a DX to a TGW or you need to have DXGW to connect a VIF to a TGW.

A is correct as a DXGW is a global resources that allows cross-region attachments

upvoted 5 times

arthurmeirelessm 5 months, 1 week ago

Por que Direct Connect Gateway nao forneceria conectividade a outras regiões?

upvoted 1 times

amministrazione 10 months ago

A. Provision a Direct Connect gateway. Delete the existing private virtual interface from the existing connection. Create the second Direct Connect connection. Create a new private virtual interface on each connection, and connect both private virtual interfaces to the Direct Connect gateway. Connect the Direct Connect gateway to the single VPC.

upvoted 1 times

MoT0ne 1 year, 3 months ago

Private Virtual Interface is a logical connection between your Direct Connect connection and a Direct Connect gateway. It is a virtual representation of the physical connection and allows you to establish connectivity to the VPCs associated with the Direct Connect gateway.

upvoted 1 times

KyleZheng 1 year, 6 months ago

A

Because “Transit GW can also communicate from on-premises to AWS, but this one uses Site to Site VPN (IPSec VPN).“

upvoted 1 times

shaaam80 1 year, 6 months ago

Selected Answer: A

Answer A. DCGW is the only option here as it supports both DC connections plus allows expansion into other regions. TGW does not span regions.
upvoted 3 times

severlight 1 year, 7 months ago

Selected Answer: A

multiple regions - dx gateway
upvoted 1 times

AMohanty 1 year, 9 months ago

None of the options seem to satisfy the condition "Solution must provide connectivity to other regions through same pair of Direct Connect Connections.

In both option A and D, we don't talk of associating second region VPC to the Transit GW or Direct Connect GW.
upvoted 1 times

whenthan 1 year, 10 months ago

Selected Answer: A

https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/
upvoted 1 times

NikkyDicky 2 years ago

Selected Answer: A

It's A.
D is not suported
upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: A

A
keyword === Direct Connect gateway
upvoted 1 times

gameoflove 2 years, 1 month ago

Selected Answer: A

A. Is the Correct Option as Direct Connect Gateway with Private Virtual Interface will meet the requirement
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: A

Provision a Direct Connect gateway.
upvoted 2 times

God_Is_Love 2 years, 4 months ago

Logical answer : B and C are good for existing architecture in question. But with redundant DX connection requirement, only solution is Gateway. that resolves to A(Direct connect gateway) or D(Transit gateway), but D as transit gateway is wrong because it mentions private interfaces connecting with transit gateway which is weird [usually VPC attachments are made connecting transit gateway]. So answer is A - Direct Connect Gateway. (Infact, this is future proof when we want different VPCs in different regions later with this architecture)
upvoted 3 times

Question #18 *Topic 1*

A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.
The website contains static content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2 instances running in an Auto Scaling group to process an Amazon SQS queue.
The company wants to re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.
Which solution meets these requirements?

A. Use Amazon ECS containers for the web application and Spot instances for the Auto Scaling group that processes the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.

B. Store the uploaded videos in Amazon EFS and mount the file system to the EC2 instances for the web application. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.

C. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notification to publish events to the SQS queue. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.

D. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue. Replace the custom software with Amazon Rekognition to categorize the videos.

Suggested Answer: *D*

Community vote distribution

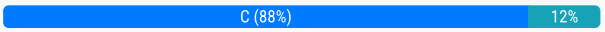

masetromain Highly Voted 9 months, 1 week ago

Selected Answer: C

This solution meets the requirements by using multiple managed services offered by AWS which can reduce the operational overhead. Hosting the web application in Amazon S3 would make it highly available, scalable and can handle variable traffic. The uploaded videos can be stored in S3 and processed using S3 event notifications that trigger a Lambda function, which calls the Amazon Rekognition API to categorize the videos. SQS can be used to process the event notifications and also it is a managed service.
This solution eliminates the need to manage EC2 instances, EBS volumes and the custom software. Additionally, using Lambda function in this case, eliminates the need for managing additional servers to process the SQS queue which will reduce operational overhead.

By using this solution, the company can benefit from the scalability, reliability, and cost-effectiveness that these services offer, which can help to reduce operational overhead and improve the overall performance and security of the application.
upvoted 32 times

Mahakali 1 year, 9 months ago

Any explanation on option A ?
upvoted 1 times

AWSum1 9 months, 3 weeks ago

ECS is managed to an extent, but the question fails to elaborate, no mention of fargate etc. There's unnecessary mentions of spot instances to confuse you. The web application has static content which can be hosted in S3 instead of ECS
upvoted 2 times

RaghavendraPrakash Highly Voted 2 years, 2 months ago

D. Because, you cannot host web application in S3, only static web assets. ElasticBeanStalk provides an easy way to onboard autoscaling web apps with minimal operational overheads.
upvoted 13 times

Kirkster 5 months, 2 weeks ago

You absolutely can host a static website in Amazon S3, I do it all the time (you create DNS records pointing to the S3 bucket), although putting CloudFront in front of it would be better. S3 web hosting even allows custom 404 error pages, and selecting a default (index.html) page, etc.
upvoted 1 times

7f6aef3 1 year, 2 months ago

Rekognition no consulta directamente EBS, pero puedes cargar datos en un recurso de almacenamiento compatible con Rekognition, como S3, para que Rekognition realice análisis sobre esos datos.

upvoted 1 times

7f6aef3 1 year, 2 months ago

Rekognition does not query EBS directly, but you can upload data to a Rekognition-compatible storage resource, such as S3, for Rekognition to perform analysis on that data.

upvoted 1 times

gofavad926 1 year, 3 months ago

"The company wants to re-architect the application "...

upvoted 2 times

Kirkster Most Recent 5 months, 2 weeks ago

Selected Answer: C

Answer A doesn't really address the operational burden, and the Spot instance stuff is a distractor. Answer C (hosting the static website in S3, along with the videos, and switching to Rekognition with S3 events) provides the most reduction in operational burden, and as a plus is also going to be the lowest-cost solution.

upvoted 1 times

Drake17 6 months, 2 weeks ago

Selected Answer: C

The Amazon Rekognition Video API facilitates the analysis of videos either stored in an Amazon S3 bucket or streamed via Amazon Kinesis Video Streams.

https://docs.aws.amazon.com/rekognition/latest/dg/how-it-works-operations-intro.html

upvoted 1 times

TariqKipkemei 7 months, 3 weeks ago

Selected Answer: C

'static website' = Amazon S3
'store videos'= Amazon S3
'video analysis' = Amazon Rekognition
'reduce operational overhead, managed service' = S3 events, AWS Lambda, Amazon SQS

upvoted 1 times

cudbyanc 9 months, 1 week ago

Selected Answer: C

The answer is C.

This solution eliminates the need for managing and scaling EC2 instances for the web application and the worker environment for processing the SQS queue. Instead, Amazon S3 can host the web application, and store the uploaded videos, which can trigger S3 event notifications to send messages to the SQS queue. Then, an AWS Lambda function can process the messages in the SQS queue and use Amazon Rekognition API to categorize the videos. This approach also takes advantage of AWS-managed services, such as S3, SQS, and Lambda, which reduces operational overhead and dependency on third-party software.

upvoted 5 times

Bereket 9 months, 1 week ago

Selected Answer: C

C. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notification to publish events to the SQS queue. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.

Explanation:
Hosting the Web Application in Amazon S3:

Cost-effective and Scalable: Amazon S3 is a cost-effective and scalable solution for hosting static web content. It can handle variable traffic efficiently without the need to manage servers.
Static Content Hosting: Ideal for serving static content like HTML, CSS, JavaScript, and media files.

upvoted 1 times

kz407 9 months, 1 week ago

Selected Answer: C

While I vote for C, I do think however that whether we can go with C really depends on the application codebase.
The use case mentions that the application enables file uploads. We know that handling files require a backend, if your application is written in something like Java. If that's the case, you won't be able to host your application in S3. The phrase "website contains static content" is really vague,

as it does not reveal anything about the backend of the application.
Now, the fact that the application has EBS to store Video files give up a hint, that suggests that the application has some BE code.
I am taking a hint from "re-architect" I assume involves some revamping of the applications codebase. So, here's how I'd go about "re-architecting"
1. Move storage of files to S3.
2. Eliminate the BE codebase, revamp the FE codebase to rely entirely on AWS JS SDK and handle file uploads with that. Now you don't need to manage any compute resources at all.
3. Go about the rest of the solution.
upvoted 1 times

924641e 9 months, 1 week ago

Selected Answer: C

The mention of static content really throws this question off and clearly the community thinks this as well. The argument of static website vs static content being the key to selecting D isn't really a strong argument but that doesn't exclude D from being a viable solution. Operational overhead is minimized with Elastic Beanstalk and removes dependencies on third party tools/software.
upvoted 2 times

24Gel 1 year, 3 months ago

thanks, this is the best explain
upvoted 1 times

grire974 9 months, 1 week ago

Selected Answer: C

If it were D - how would Rekognition access the videos to classify? Rekognition would need to ssh into the EBS volume of various beanstalk instances running under an ASG (impossible as far as I know). I agree though - I think the wording is terrible for 'contains static content'; as how on earth would this type of app practically run on s3 alone for login/ user auth etc.. would need to be coupled with other serverless products such as lambda/cognito etc.
upvoted 1 times

grire974 1 year, 5 months ago

per my previous comment; s3 is the only viable data source for rekognition
https://aws.amazon.com/rekognition/faqs/#:~:text=Amazon%20Rekognition%20Video%20operations%20can,are%20MPEG%2D4%20and%20MOV.

from my experience this is the same too with similar services like elastic transcoder
upvoted 1 times

amministrazione 10 months ago

C. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notification to publish events to the SQS queue. Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
upvoted 1 times

ff32d79 10 months, 3 weeks ago

I saw this question in other question bank (owner of the questions) and it is A, reason is assuming is moving files back and forth cannot be static page, so it is A.
upvoted 1 times

Helpnosense 1 year ago

Selected Answer: C

Only Answer C is the solution that covers all the requirements, where the videos are stored, how SQS messages are produced and consumed, how web app is hosted.
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, this is a typical scenario
upvoted 1 times

MoT0ne 1 year, 3 months ago

re-architect the application to reduce operational overhead
upvoted 1 times

subbupro 1 year, 6 months ago

Elastic bean stack is not required , it is a static content only, better can go with S3. So Answer is C
upvoted 1 times

abeb 1 year, 7 months ago

C videos in Amazon S3

upvoted 1 times

Question #19 *Topic 1*

A company has a serverless application comprised of Amazon CloudFront, Amazon API Gateway, and AWS Lambda functions. The current deployment process of the application code is to create a new version number of the Lambda function and run an AWS CLI script to update. If the new function version has errors, another CLI script reverts by deploying the previous working version of the function. The company would like to decrease the time to deploy new versions of the application logic provided by the Lambda functions, and also reduce the time to detect and revert when errors are identified.
How can this be accomplished?

A. Create and deploy nested AWS CloudFormation stacks with the parent stack consisting of the AWS CloudFront distribution and API Gateway, and the child stack containing the Lambda function. For changes to Lambda, create an AWS CloudFormation change set and deploy; if errors are triggered, revert the AWS CloudFormation change set to the previous version.

B. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code. Rollback if Amazon CloudWatch alarms are triggered.

C. Refactor the AWS CLI scripts into a single script that deploys the new Lambda version. When deployment is completed, the script tests execute. If errors are detected, revert to the previous Lambda version.

D. Create and deploy an AWS CloudFormation stack that consists of a new API Gateway endpoint that references the new Lambda version. Change the CloudFront origin to the new API Gateway endpoint, monitor errors and if detected, change the AWS CloudFront origin to the previous API Gateway endpoint.

Suggested Answer: *B*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

AWS Serverless Application Model (SAM) is a framework that helps you build, test and deploy your serverless applications. It uses CloudFormation under the hood, so it is a way to simplify the process of creating, updating, and deploying CloudFormation templates. CodeDeploy is a service that automates code deployments to any instance, including on-premises instances and Lambda functions.
With AWS SAM you can use the built-in CodeDeploy to deploy new versions of the Lambda function, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code.
You can also define CloudWatch Alarms to trigger a rollback in case of any issues.
This allows for a faster and more efficient deployment process, as well as a more reliable rollback process when errors are identified. This way you can increase the speed of deployment and reduce the time to detect and revert when errors are identified.
upvoted 29 times

calcinator423 Most Recent 1 month ago

Selected Answer: B

A and D are just obviously wrong bc they said "reduce time to deploy" and cloudformation is very very slow.

B and C are effectively the same answer, but B uses automatic, serverless, managed architecture, and C is using manual CLI scripts and manual reverting. Generally, manual anything is discouraged.
upvoted 1 times

TariqKipkemei 7 months, 3 weeks ago

Selected Answer: B

keywords:
'Code deployment, reduce deployment time, rollback, serverless' = AWS Serverless Application Model, AWS CodeDeploy
upvoted 2 times

5up3rm4n 9 months, 1 week ago

Selected Answer: B

https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html

AWS Serverless Application Model (AWS SAM) comes built-in with CodeDeploy to provide gradual AWS Lambda deployments. With just a few lines of configuration, AWS SAM does the following for you:

Deploys new versions of your Lambda function, and automatically creates aliases that point to the new version.

Gradually shifts customer traffic to the new version until you're satisfied that it's working as expected. If an update doesn't work correctly, you can roll back the changes.

Defines pre-traffic and post-traffic test functions to verify that the newly deployed code is configured correctly and that your application operates as expected.

Automatically rolls back the deployment if CloudWatch alarms are triggered.
upvoted 3 times

atirado 9 months, 1 week ago

Selected Answer: B

Option A - This work will allow reverting to previous versions of the Lambda functions but reverting means all functions will be reverted. This does not minimize the the time needed to detect and revert errors.

Option B - This option minimizes the time needed to deploy functions and detect and revert errors: As each function is deployed it can be tested and reverted individually. Moreover, the option provides a straightforward mechanism to detect and revert errors: Detect errors in CloudWatch, fix the functions' code in SAM, redeploy with AWS CodeDeploy.

Option C - This option does not minimize the time needed to detect and revert errors. It only automates the current process.

Option D - This option does not minimize the time needed to detect and revert errors: It takes time for CloudFormation to switch origins and nothing has been done to about the current process for deploying and testing functions.
upvoted 3 times

AWSum1 9 months, 3 weeks ago

Selected Answer: B

https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html

Pretty much what the question wants
upvoted 1 times

amministrazione 10 months ago

B. Use AWS SAM and built-in AWS CodeDeploy to deploy the new Lambda version, gradually shift traffic to the new version, and use pre-traffic and post-traffic test functions to verify code. Rollback if Amazon CloudWatch alarms are triggered.
upvoted 1 times

AwsZora 1 year ago

why not a
upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B, use SAM to deploy serverless applications on aws
upvoted 1 times

shaaam80 1 year, 6 months ago

Selected Answer: B

Answer B. Use SAM and Codedeploy. Revert if any errors to the previous version.
upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: B

obvious
upvoted 1 times

whenthan 1 year, 10 months ago

Selected Answer: B

requirmeents :
decrease the time to deploy new versions of the application logic provided by the Lambda functions,
revert when erros identified
upvoted 1 times

NikkyDicky 2 years ago

Selected Answer: B

B no do0ubt

upvoted 1 times

Jonalb 2 years ago

Selected Answer: B

100% B

upvoted 1 times

gameoflove 2 years, 1 month ago

Selected Answer: B

B solve the problem which is causing in the current scenario

upvoted 1 times

2aldous 2 years, 2 months ago

Selected Answer: B

Definitile B

https://docs.aws.amazon.com/es_es/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

Use AWS SAM and built-in AWS CodeDeploy

upvoted 1 times

Question #20 *Topic 1*

A company is planning to store a large number of archived documents and make the documents available to employees through the corporate intranet. Employees will access the system by connecting through a client VPN service that is attached to a VPC. The data must not be accessible to the public.
The documents that the company is storing are copies of data that is held on physical media elsewhere. The number of requests will be low. Availability and speed of retrieval are not concerns of the company.
Which solution will meet these requirements at the LOWEST cost?

A. Create an Amazon S3 bucket. Configure the S3 bucket to use the S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class as default. Configure the S3 bucket for website hosting. Create an S3 interface endpoint. Configure the S3 bucket to allow access only through that endpoint.

B. Launch an Amazon EC2 instance that runs a web server. Attach an Amazon Elastic File System (Amazon EFS) file system to store the archived data in the EFS One Zone-Infrequent Access (EFS One Zone-IA) storage class Configure the instance security groups to allow access only from private networks.

C. Launch an Amazon EC2 instance that runs a web server Attach an Amazon Elastic Block Store (Amazon EBS) volume to store the archived data. Use the Cold HDD (sc1) volume type. Configure the instance security groups to allow access only from private networks.

D. Create an Amazon S3 bucket. Configure the S3 bucket to use the S3 Glacier Deep Archive storage class as default. Configure the S3 bucket for website hosting. Create an S3 interface endpoint. Configure the S3 bucket to allow access only through that endpoint.

Suggested Answer: *D*

Community vote distribution

tman22 Highly Voted 2 years, 6 months ago

A - Glacier Deep Archive can't be used for web hosting, regardless if the company says retrieval time is no concern.

upvoted 40 times

SaqibTaqi 1 week, 1 day ago

where is it mentioned to use web hosting? lowest cost is the priority and company has no cencerns about retrieval speed? the answer would be D. i also asked ChatGPT and deepseek. both gave the same answer.

upvoted 1 times

tman22 2 years, 6 months ago

Nevermind, I go for D.
It should be technically possible - and mostly dependent on the intranet web application logic - It could present users with the ability to start file retrieval, for then to later access the data.

upvoted 18 times

zhangyu20000 Highly Voted 2 years, 6 months ago

A is correct. HA is not required here.
D use Glacier deep archive that need hours to access that will cause time out for web

upvoted 22 times

Kaps443 Most Recent 3 weeks, 2 days ago

Selected Answer: A

All components work, supports website hosting, low cost

upvoted 1 times

thiagodotoli 1 month ago

Selected Answer: D

A questão não aborda hospedagem. A intranet já existe e fará a consulta no S3. Geralmente armazenamento MENOR custo e que não tenha tempo de recuperação a resposta é o Glacier.

upvoted 1 times

calcinator423 1 month ago

Selected Answer: D

"speed of retrieval does not matter" = s3 glacier deep archive. Although S3 Glacier Deep Archive can't be used for website hosting, you also don't need it because you're accessing the documents via the S3 interface endpoint via the intranet.

upvoted 1 times

Barrieta 1 month, 2 weeks ago

Selected Answer: D

Why D is correct:

S3 Glacier Deep Archive is the lowest-cost storage class designed for rarely accessed data, which fits the use case perfectly.

Access through an S3 interface endpoint ensures that data is not exposed to the public internet, meeting the security requirement.

Hosting the files through S3 website hosting (even though retrieval is slow) provides a simple mechanism for access via intranet (assuming internal DNS can resolve the endpoint).

There's no need to run or manage EC2 instances, no web servers, and no block or file storage, reducing both operational overhead and cost.

upvoted 1 times

ausl 1 month, 3 weeks ago

Selected Answer: A

A - Glacier Deep Archive can't be used for web hosting, regardless if the company says retrieval time is no concern.

upvoted 1 times

proawsk 2 months ago

Selected Answer: D

Why are people bringing in webhosting here, it said it would be accessible over corporate intranet. So their they can implement to raise a request for accessing any archive and once it's ready they can share a link that can be accessed over the intranet. So D is the right answer.

upvoted 2 times

BennyMao 3 months, 3 weeks ago

Selected Answer: D

S3 Glacier Deep Archive will be the lowest cost.

upvoted 1 times

eberhe900 4 months, 1 week ago

Selected Answer: D

The question does not require website hosting. Its focus is on securely storing documents so employees can access them via VPN without exposing them to the public Internet. So Glacier Deep Archive is more cost-saving

upvoted 3 times

ahhatem 6 months, 3 weeks ago

Selected Answer: D

Buckets do not have a storage class... objects do! You can enable web hosting on any bucket and retrieve the objects (when available) over HTTP, whether the objects are available instantly or not that is another question. For glacier, you will have to request a restore and wait but when the object is eventually ready, you can download it directly over HTTP.

So, the answer is D.

upvoted 3 times

Heman31in 6 months, 3 weeks ago

Selected Answer: A

S3 glacier can be used for hosting as it requires additional steps to unarchive objects!

upvoted 1 times

wem 6 months, 3 weeks ago

Selected Answer: A

No, Amazon S3 Glacier Deep Archive cannot be used for web hosting.

Reasons:

Glacier Deep Archive is not designed for real-time access:

Glacier Deep Archive is intended for long-term, infrequently accessed archival data. Retrieval times can range from 12 hours to 48 hours, making it unsuitable for serving content dynamically or in real-time through a website.

No direct web hosting capabilities:

S3 web hosting requires immediate availability of objects stored in the bucket. S3 Glacier Deep Archive is designed for delayed access, so objects stored in this class cannot be served directly.
Storage class retrieval process:

Objects stored in Glacier or Glacier Deep Archive require a retrieval job to be initiated before they are available for access. This process is asynchronous and incompatible with the demands of a web-hosting use case.
upvoted 1 times

TariqKipkemei 7 months, 3 weeks ago

Selected Answer: D

keywords:
'Archive documents, low requests, low availability and speed, LOWEST cost' = S3 Glacier Deep Archive
upvoted 1 times

TewatiaAmit 8 months, 1 week ago

Selected Answer: A

only A & D make sense, and since website hosting is also a requirement, I'll go with A.
upvoted 1 times

c73bf38 9 months, 1 week ago

Selected Answer: A

The requirements are to store a large number of archived documents that are not publicly accessible, and make them available to employees through a corporate intranet. As the number of requests is low and speed of retrieval is not a concern, we can use the low-cost S3 One Zone-Infrequent Access (S3 One Zone-IA) storage class. We can configure the S3 bucket for website hosting and create an S3 interface endpoint to allow access to the documents only through the corporate intranet. This solution is the lowest cost as it eliminates the need to launch and manage EC2 instances.

Option B and C involve launching an EC2 instance which increases the operational overhead and is more expensive than using S3. Also, EFS One Zone-IA storage class is not recommended for storing large files.

Option D involves using the S3 Glacier Deep Archive storage class which is intended for long-term archival storage of data and not suitable for retrieving data frequently.
upvoted 4 times

MRL110 1 year, 11 months ago

S3 interface endpoint doesn't support web hosting.
The question does not say large files, but large number of archived documents, which could be small-sized. Hence EFS OZ-IA (being cheaper than SC1) could be the right answer.
upvoted 1 times

zejou1 9 months, 1 week ago

Selected Answer: A

https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-class-intro.html
Store large number of archived docs, and available through corp intranet.
Copies of data held on physical media elsewhere (could be re-created).
Requests low (but it doesn't say RARE so think monthly/quarterly).
"AVAILABILITY" and speed of retrieval are not concerns.

It is A, yes Glacier is "cheaper", but I have to leave the archives for at least 180 days, would be available on corp intranet and it is more cost-effective if I want to migrate the data to Glacier if I monitor use and see it is "rarely" touched and know I have to hold it due to regulatory for at minimal 180 days.
upvoted 4 times

Question #21 *Topic 1*

A company is using an on-premises Active Directory service for user authentication. The company wants to use the same authentication service to sign in to the company's AWS accounts, which are using AWS Organizations. AWS Site-to-Site VPN connectivity already exists between the on-premises environment and all the company's AWS accounts.
The company's security policy requires conditional access to the accounts based on user groups and roles. User identities must be managed in a single location.
Which solution will meet these requirements?

A. Configure AWS IAM Identity Center (AWS Single Sign-On) to connect to Active Directory by using SAML 2.0. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. Grant access to the AWS accounts by using attribute-based access controls (ABACs).

B. Configure AWS IAM Identity Center (AWS Single Sign-On) by using IAM Identity Center as an identity source. Enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. Grant access to the AWS accounts by using IAM Identity Center permission sets.

C. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use a SAML 2.0 identity provider. Provision IAM users that are mapped to the federated users. Grant access that corresponds to appropriate groups in Active Directory. Grant access to the required AWS accounts by using cross-account IAM users.

D. In one of the company's AWS accounts, configure AWS Identity and Access Management (IAM) to use an OpenID Connect (OIDC) identity provider. Provision IAM roles that grant access to the AWS account for the federated users that correspond to appropriate groups in Active Directory. Grant access to the required AWS accounts by using cross-account IAM roles.

Suggested Answer: *D*

Community vote distribution

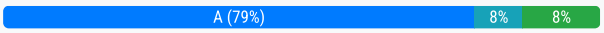

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

https://www.examtopics.com/discussions/amazon/view/74174-exam-aws-certified-solutions-architect-professional-topic-1/

Both option C and option A are valid solutions that meet the requirements for the scenario.

ABAC, or attribute-based access control, is a method of granting access to resources based on the attributes of the user, the resource, and the action. This allows for fine-grained access control, which can be useful for implementing a security policy that requires conditional access to the accounts based on user groups and roles.

AWS IAM Identity Center (AWS SSO) allows you to connect to your on-premises Active Directory service using SAML 2.0. With this, you can enable automatic provisioning by using the System for Cross-domain Identity Management (SCIM) v2.0 protocol, which allows for the management of user identities in a single location.
upvoted 30 times

masetromain 2 years, 5 months ago

In option C, the company will use IAM to use a SAML 2.0 identity provider, and it will use the appropriate groups in Active Directory to grant access to the required AWS accounts by using cross-account IAM users. In this way, it can implement its security policy of conditional access to the accounts based on user groups and roles.

In summary, both option A and C are valid solutions, both of them allow you to use your on-premises Active Directory service for user authentication, and both of them allow you to manage user identities in a single location and grant access to the AWS accounts based on user groups and roles.
upvoted 2 times

bititan Highly Voted 2 years, 5 months ago

Selected Answer: A

A is has options for SAML and SCIM configuration with AD
C is all about users and no roles are mentioned. AD User attributes cannot be mapped to IAM users direct
D is openID based, MS AD would not support this

so I go with A
upvoted 14 times

trap 1 year, 7 months ago
native AD doesn't support SAML 2.0 without an ADFS server. SCIM is also not supported at all. SCIM provisioning is supported by other IDPs like Azure AD
upvoted 3 times

gonzjo52 1 year, 2 months ago
Si, si son compatibles. https://aws.amazon.com/es/directoryservice/faqs/
upvoted 1 times

trap 1 year, 7 months ago
https://docs.aws.amazon.com/singlesignon/latest/userguide/supported-idps.html
upvoted 2 times

Heman31in Most Recent 6 months, 3 weeks ago
Selected Answer: A
SCIM v2.0 in the Context of AWS SSO and Active Directory
When using AWS IAM Identity Center (AWS SSO) with Active Directory, SCIM v2.0 is utilized to automatically provision and de-provision users and groups. This eliminates the need for manual user or group management and ensures that changes in your on-premises AD are reflected in AWS SSO.
upvoted 1 times

TariqKipkemei 7 months, 3 weeks ago
Selected Answer: A
keywords:
'Use active directory to sign in Aws accounts' = AWS IAM Identity Center (AWS Single Sign-On) , SAML 2.0
'conditional access to the accounts based on user groups and roles' = AWS IAM (ABAC)
upvoted 2 times

Ashu_0007 10 months, 2 weeks ago
AWS IAM Identity Center + SAML
upvoted 1 times

Vaibs099 1 year, 5 months ago
A is correct

Reasons -

Option A mentions about Active Directory as identity Source configuration which solves the purpose of establishing trust and sync from on prem AD using Directory Service. Solves the purpose of using on-prem AD as Single Sign On asked in the question.

It is also mentioned that AWS org is in place, which works well with AWS Identity Centre. Gives another validation. It gives us hint of efficiently managing AWS Org accounts / OUs with Identity Centre (Permission Set behind the scene) to manage RBAC within accounts.

Finally this line - "The company's security policy requires conditional access to the accounts based on user groups and roles." is talking about conditional access which can only be solved by ABAC(Attribute Based Access Control). For example user with green attribute should only get access to resources with green attribute. This can be solved by Tag functionality within AWS Identity Centre.
upvoted 3 times

atirado 1 year, 6 months ago
Selected Answer: D
Option A - This option works however it moves authentication and managing user identities from Active Directory to Identity Center but the question states the company wants to use the same authentication service to sign into AWS in reference to Active Directory

Option B - This option works but it moves user identity management and authentication tie Identity Center which is not what the question states the company wants to do

Option C - This option does not work because in AWS you provision cross-account IAM roles rather than users.

Option D - This option might work but it is missing AD FS, a component that enables OIDC flows in AD. Otherwise it maintains user identity management in one place and allows the company to keep using Active Directory for authentication as the question states

upvoted 2 times

ninomfr64 1 year, 6 months ago

Selected Answer: B

Didn't spent time checking if C and D works, because when you have an AWS Organitazion and need to use AD to sign-in to the company's AWS accounts AWS IdC is the way to go.

Now, with AWS IdC we need ADFS and while ADFS does not support SCIM, it is possible to still have your users and groups automatically synchronize with the IAM IDC by using the SCIM API and PowerShell as per https://aws.amazon.com/blogs/modernizing-with-aws/synchronize-active-directory-users-to-aws-iam-identity-center-using-scim-and-powershell/#:~:text=While%20ADFS%20does%20not%20support,the%20SCIM%20API%20and%20PowerShell.

Finally, ABAC is an authorization strategy and it is not alternative to IdC Permission Sets. Also the scenario requires conditional access to the accounts based on user groups and roles, this point me to RBAC strategy. I would pick ABAC if the request mentioned user attributes like Department, Cost Center or Project thus.

upvoted 2 times

ninomfr64 1 year, 5 months ago

After reviewing it, the correct answer is A. "User identities must be managed in a single location" -> "Configure AWS IAM Identity Center (AWS Single Sign-On) to connect to Active Directory by using SAML 2.0" while B states "Configure AWS IAM Identity Center (AWS Single Sign-On) by using IAM Identity Center as an identity source". Using AWs IdC as identity source will not meet requirement to manage all users in a single place

upvoted 1 times

924641e 1 year, 6 months ago

Answer A for AWS SSO would the right answer at first glance since IAM roles can be mapped to AD groups but it would require additional AD functions like ADFS for SCIM so the next best option is D.

upvoted 3 times

subbupro 1 year, 6 months ago

A is a correct one, because need to use the SAML for single sign on from the on-premise directory and also C is not correct because the federated should not come in to the picture federated is for only facebook, twitter, gmail account sign on - but we should use the companies active directory, so A is a correct one.

upvoted 1 times

siasiasia 1 year, 7 months ago

Selected Answer: C

AD and SCIM don't go together so forget A and B. I've never seen a document talking about integrating OpenID with AWS account login so D is also out. C is doable so I go with C.

upvoted 1 times

gonzjo52 1 year, 2 months ago

P: ¿Puedo usar la autenticación basada en lenguaje de marcado de aserción de seguridad (SAML) 2.0 con aplicaciones de la nube que usen AWS Managed Microsoft AD?

Sí. Puede usar los servicios federados de Microsoft Active Directory (AD FS) para Windows 2016 con su dominio administrado de AWS Managed Microsoft AD para autenticar usuarios en aplicaciones en la nube compatibles con SAML. https://aws.amazon.com/es/directoryservice/faqs/

upvoted 1 times

sizzla83 1 year, 7 months ago

I am with B on this one. A is incorrect because you can only use ABAC (Attribute-Based Access Control) with IAM Identity Center Identity Store NOT with Active Directory

upvoted 1 times

ninomfr64 1 year, 6 months ago

Agree with you on B, but:

- You can use IAM Identity Center to manage access to your AWS resources across multiple AWS accounts using user attributes that come from any IAM Identity Center identity source - https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html

- ABAC is an authorization strategy that defines permissions based on attributes and it is implemented using IdC Permission Sets.

upvoted 1 times

enk 1 year, 7 months ago

Selected Answer: A

As mentioned, SAML 2.0 doesn't directly integrate with AD and requires ADFS proxy as a go between, so the lack of ADFS being mentioned in A or B is throwing people off. However, AD on-premise with direct/VPN connectivity...IAM identify center is the way to go for SSO. I believe ADFS is implied when the question casually mentions "IAM Identify Center connect to AD using SAML 2.0".

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: A

federated IdP is required and access to multiple accounts

upvoted 1 times

trap 1 year, 7 months ago

Answer A and B are wrong!!!

Active Directory doesn't support SAML without the use of Active Directory Federation Server!! SCIM is also not supported. The articles that all are pasting here mention the need of an AD connect or the trust between the local AD and an AWS managed Microsoft AD which is not the case here.

C is also wrong. Cross account IAM users option doesn't exist.

The correct is D!! You can use an OpenID Connect (OIDC) identity provider (e.g OKTA or Azure AD) and sync AD groups in it. You can then use cross account roles to grant access to the federated users

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_create_oidc.html

https://help.okta.com/en-us/content/topics/directory/ad-agent-manage-users-groups.htm

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_aws-accounts.html

upvoted 3 times

M4D3V1L 1 year, 9 months ago

Selected Answer: A

https://docs.aws.amazon.com/singlesignon/latest/userguide/onelogin-idp.html#onelogin-passing-abac

upvoted 1 times

imvb88 1 year, 9 months ago

Selected Answer: A

A: combination SSO + SAML2.0 + AD sounds correct. Automatic provisioning with SCIM means creating users and groups that synced with AD. ABAC seems not too fit for this as the requirements is "requires conditional access to the accounts based on user groups and roles" but that already satisfied with SCIM.

B: "use Identity Center as an identity source" -> not using on premise AD -> wrong

D: use OIDC -> wrong as on premise AD does not support OIDC. Cannot find an exact source for this but ChatGpt says so..

C: creating users mapped to federated users sounds red flags. Could have been correct if it was "creating roles", the same way with the classic "creating roles for EC2 to access S3 instead of user..."

Conclusion: A

upvoted 3 times

Question #22 *Topic 1*

A software company has deployed an application that consumes a REST API by using Amazon API Gateway, AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

A. Implement retry logic with exponential backoff and irregular variation in the client application. Ensure that the errors are caught and handled with descriptive error messages.

B. Implement API throttling through a usage plan at the API Gateway level. Ensure that the client application handles code 429 replies without error.

C. Turn on API caching to enhance responsiveness for the production stage. Run 10-minute load tests. Verify that the cache capacity is appropriate for the workload.

D. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.

Suggested Answer: *B*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

API throttling is a technique that can be used to control the rate of requests to an API. This can be useful in situations where a small number of clients are making a large number of requests, which is causing errors. By implementing API throttling through a usage plan at the API Gateway level, the solutions architect can limit the number of requests that a client can make, which will help to reduce the number of errors.

It's important that the client application handles the code 429 replies without error, this will help to improve the customer experience by reducing the number of errors that are displayed to customers. Additionally, it will prevent the API's reputation from being damaged by the errors.
upvoted 45 times

masetromain 2 years, 5 months ago

It is important to note that other solutions such as retry logic with exponential backoff and irregular variation in the client application or turn on API caching to enhance responsiveness for the production stage may help to improve the customer experience and reduce errors, but they do not address the root cause of the problem which is a large number of requests coming from a small number of clients.

Implementing reserved concurrency at the Lambda function level can provide resources that are needed during sudden increases in traffic, but it does not address the issue of a client making a large number of requests and causing errors.
upvoted 16 times

zhangyu20000 Highly Voted 2 years, 6 months ago

B is correct. API gateway throttling is applied to single account - https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html. Retry will make it even worse.
upvoted 8 times

HajaMydin Most Recent 1 month ago

Selected Answer: B

Why B is the best solution:
Problem Summary:
Increased errors during PUT requests.

Caused by a single client generating a high volume of requests.

The API is noncritical, and retries are acceptable.

Errors are affecting customer experience and API reputation.

Solution Breakdown:
API Gateway usage plans allow you to throttle requests by API key, limiting the impact of a noisy or abusive client.

Throttling results in HTTP 429 (Too Many Requests) responses.

Clients can be coded to recognize 429 responses and retry with backoff, which improves customer experience while maintaining control.

This approach protects backend systems like Lambda and DynamoDB from being overloaded.
upvoted 1 times

f3f4935 2 months, 1 week ago

Selected Answer: B

B is correct
upvoted 1 times

TariqKipkemei 7 months, 1 week ago

Selected Answer: B

keywords:
'A large number of the PUT requests' = API throttling
upvoted 1 times

amministrazione 10 months ago

B. Implement API throttling through a usage plan at the API Gateway level. Ensure that the client application handles code 429 replies without error.
upvoted 1 times

Ashu_0007 10 months, 2 weeks ago

API gateway throttling
upvoted 1 times

Jason666888 11 months ago

Selected Answer: B

Key word: a large number of the PUT requests, one client

Seeing this will ring a bell on throttling on API Gateway. But normally you also need to make sure when the client side see "429 too many attempts", the app can capture that error code and show some reasonable error message(e.g. You have sent too many requests .Please try again later)
upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B. C only will help with GET requests, and A and D don't prevent it
upvoted 1 times

anubha.agrahari 1 year, 3 months ago

Selected Answer: B

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html
upvoted 1 times

duriselvan 1 year, 4 months ago

B ans : https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-request-throttling.html
upvoted 1 times

AimarLeo 1 year, 4 months ago

This question missing MASSIVE information.. none of the answers can fulfil the requirements..
upvoted 2 times

bjexamprep 1 year, 5 months ago

Selected Answer: A

There is no evidence indicating the problem is with the throughput. If it is throughput, other clients will have similar problem.
And “the errors are displayed to customers and are causing damage to the API’s reputation.”, this means the solution should be able to reduce the error message showed on the client side, while, throttling the client will actually close the service for this particular client, which is against the “clients can tolerate retries of unsuccessful calls”.
I vote A for this question.
upvoted 1 times

sarfaraz_khan 1 year, 6 months ago

The solutions architect should recommend option B: Implement API throttling through a usage plan at the API Gateway level. Ensure that the client application handles code 429 replies without error.

Option B is the most directly related recommendation to improving the customer experience, as it addresses the issue of API rate limiting and ensures a more predictable and controlled experience for users.

upvoted 1 times

atirado 1 year, 6 months ago

Selected Answer: B

Option A - This option will make retries take longer on each retry for all clients rather than for the few causing issues in the application

Option B - This option will work: An usage plan will allow throttling requests from specific clients identified by their API Key and ensuring client applications can handle throttling errors provides a consistent experience

Option C - This option has no relation with the problem at hand

Option D - This option assumes there is a capacity issue managing the increase in volumes but given that errors occur due to a small number of clients then reserved concurrency will not address the cause of the issue

upvoted 2 times

atirado 1 year, 6 months ago

Selected Answer: B

Option A - This option will make retries take longer on each retry for all clients rather than for the few causing issues in the application

Option B - This option will work: An usage plan will allow throttling requests from specific clients identified by their API Key and ensuring client applications can handle throttling errors provides a consistent experience

Option C - This option has no relation with the problem at hand

Option D - This option assumes there is a capacity issue managing the increase in volumes but given that errors occur due to a small number of clients then reserved concurrency will not address the cause of the issue

upvoted 1 times

ninomfr64 1 year, 6 months ago

Selected Answer: B

Usage Plan throttling prevents a group of users or a single user to saturate the API concurrency capacity. Thus B. Also A and D can help in this scenario, but they will have less benefit with respect to B. While C does not help in this scenario as I do not see how API Gateway caching can help PUT requests

upvoted 1 times

Question #23 *Topic 1*

A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.
A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.
Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.

B. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template. Use the EBS volume as the shared storage for the duration of the job. Detach the EBS volume when the job is complete

C. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.

D. Migrate the data from the existing shared file system to an Amazon S3 bucket. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job. Delete the file gateway when the job is complete.

Suggested Answer: *D*

Community vote distribution

sambb Highly Voted 2 years, 4 months ago

Selected Answer: A

A: Lazy loading is cost-effective because only a subset of data is used at every job
B: There are hundreds of EC2 instances using the volume which is not possible (one EBS volume is limited to 16 nitro instances attached)
C: Batching would load too much data
D: storage gateway is used for on premises data access, I don't know is you can install a gateway in AWS, but Amazon would never advise this
upvoted 21 times

b3llman 1 year, 10 months ago
file storage gateway can be installed on EC2 and it is exactly used for accessing S3 from EC2 as a file system
upvoted 1 times

Chainshark 1 year, 8 months ago
It's used a lot, I've used it for customers to access and analyze data imported via Snowball from Windows machines.
upvoted 1 times

dqwsmwwvtgxwkvgcvc 1 year, 10 months ago
There is one S3 file gateway

https://aws.amazon.com/storagegateway/file/s3/
upvoted 1 times

Tofu13 1 year, 9 months ago
https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and-amazon-s3/
upvoted 3 times

chico2023 Highly Voted 1 year, 10 months ago
Answer: D

I think the main point here is to understand what they mean by "The file system must provide high performance access to the needed data" while

"provide the LARGEST overall cost reduction"?

For answer A, we have to remember that lazy load is SLOW for the first time you try to access the file (as it is being fetched from S3), BUT, as we are talking about hundreds of instances, then it might be OK. S3 Intelligent-Tiering, although doesn't seem to fit much, the part that says "The job runs once monthly, reads a subset of the files from the shared file system", indicates that at least part of the 200TB of data won't be accessed, which helps not going for answer C, for example.

My only issue with answer D is that Storage Gateway can be slower than FSx for Lustre, HOWEVER, what is the cost X performance ratio they are seeking here? We can guess that costs trumps maximum performance here: "Which solution will provide the LARGEST overall cost reduction" and, as Storage Gateway is way cheaper than FSx for Lustre per TB, it's safe to say that D is the most correct answer.

upvoted 15 times

diazed Most Recent 2 months, 2 weeks ago

Selected Answer: A

With our S3 objects imported into our Lustre file system, we can now lazy-load the files we need by simply reading the particular files. After a file is lazy-loaded, its contents are fully copied from S3 onto the Amazon FSx for Lustre file system, where it can be accessed with extremely low latency. I will go for A. https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and-amazon-s3/

upvoted 1 times

SIJUTHOMASP 6 months ago

Selected Answer: D

I lean more towards D but I am not sure whether the Gateway is only intended for on-premise as few are mentioning here. If that is not the case then the right option is D.

upvoted 1 times

zaxxon 6 months, 3 weeks ago

Selected Answer: D

FSx for Lustre, is only for Linux where in the question is Linux noted. It's states only EC2 instance not which OS is on it!

upvoted 1 times

TariqKipkemei 7 months, 1 week ago

Selected Answer: A

'The job runs once monthly', 'cost reduction' = S3 Intelligent-Tiering storage class, lazy loading.
'Scalable file system', 'shared file system', 'data-intensive ' = Amazon FSx for Lustre

upvoted 1 times

0b43291 7 months, 2 weeks ago

Selected Answer: A

By choosing Option A, the company can leverage the cost-effectiveness of Amazon S3 Intelligent-Tiering for storage and the high performance of Amazon FSx for Lustre for temporary file access, while minimizing the overall cost by creating and deleting the file system only when needed.
Option B (using Amazon EBS Multi-Attach) is not ideal because EBS volumes are designed for persistent storage, and attaching and detaching a large volume to multiple instances can be time-consuming and potentially disruptive.
Option C (using Amazon FSx for Lustre with batch loading) is less cost-effective than Option A because batch loading requires loading the entire 200 TB of data into the file system, which can be expensive and time-consuming.
Option D (using AWS Storage Gateway File Gateway) is not the most cost-effective solution because the File Gateway is designed for on-premises file storage integration and may not provide the same level of performance as FSx for Lustre for this data-intensive workload.

upvoted 2 times

amministrazione 10 months ago

A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

A or D : confusion. I wish they can provide explanation about their answers when it is not the most voted one

upvoted 1 times

Helpnosense 1 year ago

I vote D instead A because the requirement in the question is "modifies the data on the shared file system" Fsx imported data from s3 and lost the relationship to s3 after import is done Without explicitly copy back to s3, the change stays on shared file system only. Answer A solution doesn't provide a step to copy the modification back to s3.
Storage gateway presents s3 storage to the OS as shared file system. Any modification on the shared file system will be automatically saved on s3.

upvoted 3 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A: Lazy loading is cost-effective because only a subset of data is used at every job

upvoted 1 times

kz407 1 year, 3 months ago

Selected Answer: A

Problem with D is that, AWS Storage GW and File GW are solutions for integrating on-premise storage with AWS storage solutions, particularly (but not limited to) S3.

https://aws.amazon.com/storagegateway/
https://aws.amazon.com/storagegateway/file

Compute resources are residing in AWS, so having Storage GW and File GW won't solve a thing.
As far as option B is concerned, it comes down to the limitations of EBS (such as the max block size, and max number of instance that can be attached etc). Also, attaching and detaching of the EBS volumes seems a bit complicated too. On top of that, EBS does not offer the cost optimizations offered by S3 Intelligent Tiering. The question clearly mentions that only a subset of the data will be used. Intelligent tiering ensures a substantial cost optimization over time.
Hence, the answer should be A.

upvoted 3 times

kspendli 1 year, 3 months ago

Option D, migrating the data to an Amazon S3 bucket and using AWS Storage Gateway, seems to provide the largest overall cost reduction while meeting the requirements of high-performance access during the job run and minimizing costs when the storage is not actively being used. Therefore, Option D is the most suitable choice.

upvoted 1 times

anubha.agrahari 1 year, 3 months ago

Selected Answer: A

https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and-amazon-s3/

upvoted 2 times

atirado 1 year, 6 months ago

Selected Answer: A

Option A - This option might work. However, AWS FSx for Lustre does not have a feature called "lazy loading" - its default behavior is to load a file from S3 when it is first accessed (restore). It can provide high-performance as needed though nothing is said in the question about whether a slow initial load time due to restore operations could be an issue. S3 Intelligent-Tiering minimizes storage costs.

Option B - This option will provide a high-performance storage option. However, storage in EBS is expensive compared to other AWS storage services

Option C - This option might work. However, AWS FSx for Luster does not have a feature called "batch loading". Files can be pre-loaded issuing a hsm-restore command. S3 Standard is a cheap storage option yet not the cheapest option in S3

Option D - This option does not work as described in the option

upvoted 2 times

AimarLeo 1 year, 5 months ago

Actually AWS FSx for Lustre does not have a direct feature 'Lazy loading' but the question is the support of that when Amazon FSx will import the objects in our S3 bucket as files, and “lazy-load” the file contents from S3 when first access the files.. Any data processing job on Lustre with S3 as an input data source can be started without Lustre doing a full download of the dataset
first - Data is lazy loaded: only the data that is
actually processed is loaded, meaning you can
decrease your costs and latency

upvoted 1 times

ninomfr64 1 year, 6 months ago

Not B because using EBS still involves EC2 instances that are expensive (the instances that host the shared file system run continuously). Also, multi-attach is supported only for io1/oi2 EBS disk types that are expensive;
Not C as batch loading does not exists in the doc/console, I think they might refer to the option to pre-populate the data using lfs hsm_restore command as mentioned here https://docs.aws.amazon.com/fsx/latest/LustreGuide/preload-file-contents-hsm-dra.html. This would be a more

expensive option

Not D as Storage Gateway provide less performance than FSx for Lustre and it requires at least an EC2 instance and this will introduce additional cost

AA is a viable option as S3 is cheaper storage, FSx for Lustre provides performance. Lazy loading allows to actually move in the filesystem data that is actually used and intelligent tiering make sure those files that are not used are moved to less expensive S3 storage tiers.

upvoted 1 times

subbupro 1 year, 6 months ago

Intelligent tiering is not required, because the job would be running for every month, so there is no purpose for intelligent tiering, The question is having cost impact also one of the option. So go with option D.

upvoted 1 times

e4bc18e 1 year, 3 months ago

"Only a subset of data is accessed each run" So that means after 30 days data can tier down so yes there is cost savings in using INT

upvoted 1 times

Question #24 *Topic 1*

A company is developing a new service that will be accessed using TCP on a static port. A solutions architect must ensure that the service is highly available, has redundancy across Availability Zones, and is accessible using the DNS name my.service.com, which is publicly accessible. The service must use fixed address assignments so other companies can add the addresses to their allow lists.
Assuming that resources are deployed in multiple Availability Zones in a single Region, which solution will meet these requirements?

A. Create Amazon EC2 instances with an Elastic IP address for each instance. Create a Network Load Balancer (NLB) and expose the static TCP port. Register EC2 instances with the NLB. Create a new name server record set named my.service.com, and assign the Elastic IP addresses of the EC2 instances to the record set. Provide the Elastic IP addresses of the EC2 instances to the other companies to add to their allow lists.

B. Create an Amazon ECS cluster and a service definition for the application. Create and assign public IP addresses for the ECS cluster. Create a Network Load Balancer (NLB) and expose the TCP port. Create a target group and assign the ECS cluster name to the NLCreate a new A record set named my.service.com, and assign the public IP addresses of the ECS cluster to the record set. Provide the public IP addresses of the ECS cluster to the other companies to add to their allow lists.

C. Create Amazon EC2 instances for the service. Create one Elastic IP address for each Availability Zone. Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set.

D. Create an Amazon ECS cluster and a service definition for the application. Create and assign public IP address for each host in the cluster. Create an Application Load Balancer (ALB) and expose the static TCP port. Create a target group and assign the ECS service definition name to the ALB. Create a new CNAME record set and associate the public IP addresses to the record set. Provide the Elastic IP addresses of the Amazon EC2 instances to the other companies to add to their allow lists.

Suggested Answer: *C*

Community vote distribution

C (100%)

God_Is_Love Highly Voted 2 years, 4 months ago

Logical answer : Non http port like TCP should hint to NLB immediately.(ALB does not fit here) Sharing IP address of EC2 is not apt whether it is from individual EC2 instances or those from ECS cluster.this eliminates A,B.D, infact the NLB's address which stays in front of / associates to ec2 instances need to be shared. So, only solution is C

upvoted 12 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

A more appropriate solution would be option C. Create an Amazon EC2 instances for the service. Create one Elastic IP address for each Availability Zone. Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set. As it uses the NLB as the resource in the A-record, traffic will be routed through the NLB, and it will automatically route the traffic to the healthy instances based on the health checks and also it provides the fixed address assignments as the other companies can add the NLB's Elastic IP addresses to their allow lists.

upvoted 6 times

TariqKipkemei Most Recent 7 months, 1 week ago

keywords:
'TCP', 'fixed address', 'regional' = NLB

upvoted 1 times

TariqKipkemei 7 months, 1 week ago

Answer is C

upvoted 1 times

0b43291 7 months, 2 weeks ago

Selected Answer: C

By choosing Option C, the company can meet the requirements of high availability, redundancy across Availability Zones, accessibility through a DNS name (my.service.com), and fixed IP address assignments that can be added to allow lists by other companies.

Option A is not suitable because it involves creating Elastic IP addresses for each EC2 instance, which can become difficult to manage and does not provide the desired DNS name accessibility.
Option B is not appropriate because it uses an Amazon ECS cluster with public IP addresses, which may not provide the desired fixed IP addresses for allow listing by other companies.
Option D is not the correct choice because it uses an Application Load Balancer (ALB), which is designed for HTTP/HTTPS traffic and may not be the best fit for a TCP-based service. Additionally, it involves creating public IP addresses for each host in the ECS cluster, which can be complex and may not provide the desired fixed IP addresses for allow listing.
upvoted 1 times

amministrazione 10 months ago
C. Create Amazon EC2 instances for the service. Create one Elastic IP address for each Availability Zone. Create a Network Load Balancer (NLB) and expose the assigned TCP port. Assign the Elastic IP addresses to the NLB for each Availability Zone. Create a target group and register the EC2 instances with the NLB. Create a new A (alias) record set named my.service.com, and assign the NLB DNS name to the record set.
upvoted 1 times

Ashu_0007 10 months, 2 weeks ago
Ec2+NLB
upvoted 1 times

Alawi_Amazon_AWS 1 year, 2 months ago
A looks ok
https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/Strategies.html
upvoted 1 times

gofavad926 1 year, 3 months ago
Selected Answer: C
C: NLB with elastic IPs
upvoted 1 times

Vaibs099 1 year, 5 months ago
C is the right answer - Few key points - TCP static Port (go with NLB), IP Whitelisting required which can only be done with NLB. ALB doesn't support static IPs. And sharing Static (Elastic) IPs of instances of no use when using NLB. We need to share NLB Elsatic IPs from Multi AZs and create DNS record for NLB Domain Name to Domain entry.
upvoted 1 times

sammyhaj 1 year, 5 months ago
https://repost.aws/knowledge-center/elb-attach-elastic-ip-to-public-nlb
upvoted 1 times

Simon523 1 year, 9 months ago
Selected Answer: C
Other option haven't mention multi AZ
upvoted 1 times

Christina666 1 year, 11 months ago
Selected Answer: C
Static IP-> NLB
upvoted 1 times

NikkyDicky 2 years ago
Selected Answer: C
I suppose C, although you can;'t do this with A record, only alias
upvoted 1 times

SkyZeroZx 2 years ago
Selected Answer: C
Create one Elastic IP address for each Availability Zone.
upvoted 2 times

AWS_Sam 2 years, 1 month ago
C is the only option that talks about more than one AZ.
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

Create Amazon EC2 instances for the service. Create one Elastic IP address for each Availability Zone.

upvoted 2 times

kiran15789 2 years, 3 months ago

Selected Answer: C

IP address using NLB

upvoted 1 times

Question #25 *Topic 1*

A company uses an on-premises data analytics platform. The system is highly available in a fully redundant configuration across 12 servers in the company's data center.
The system runs scheduled jobs, both hourly and daily, in addition to one-time requests from users. Scheduled jobs can take between 20 minutes and 2 hours to finish running and have tight SLAs. The scheduled jobs account for 65% of the system usage. User jobs typically finish running in less than 5 minutes and have no SLA. The user jobs account for 35% of system usage. During system failures, scheduled jobs must continue to meet SLAs. However, user jobs can be delayed.
A solutions architect needs to move the system to Amazon EC2 instances and adopt a consumption-based model to reduce costs with no long-term commitments. The solution must maintain high availability and must not affect the SLAs.
Which solution will meet these requirements MOST cost-effectively?

A. Split the 12 instances across two Availability Zones in the chosen AWS Region. Run two instances in each Availability Zone as On-Demand Instances with Capacity Reservations. Run four instances in each Availability Zone as Spot Instances.

B. Split the 12 instances across three Availability Zones in the chosen AWS Region. In one of the Availability Zones, run all four instances as On-Demand Instances with Capacity Reservations. Run the remaining instances as Spot Instances.

C. Split the 12 instances across three Availability Zones in the chosen AWS Region. Run two instances in each Availability Zone as On-Demand Instances with a Savings Plan. Run two instances in each Availability Zone as Spot Instances.

D. Split the 12 instances across three Availability Zones in the chosen AWS Region. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservations. Run one instance in each Availability Zone as a Spot Instance.

Suggested Answer: *C*

Community vote distribution

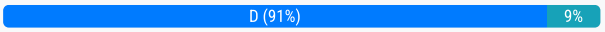

lasco Highly Voted 2 years, 4 months ago

Selected Answer: D

Voted D because of the 65% / 35% proportion. C seems to be good but with only 50% instances available we break the SLA

upvoted 25 times

joefromnc Highly Voted 1 year, 10 months ago

Can not be C because Savings Plans requirement long term commitment.

upvoted 7 times

BennyMao Most Recent 3 months, 3 weeks ago

Selected Answer: C

Savings Plans provide cost savings compared to On-Demand while ensuring predictable compute resources for scheduled jobs with tight SLAs.
Balanced distribution of On-Demand and Spot Instances across AZs ensures redundancy and cost-effectiveness.

upvoted 1 times

TariqKipkemei 7 months, 1 week ago

Selected Answer: D

keywords:
'65%, more than half of the instances must continue to meet SLAs' = On-demand instances with capacity reservations.
'cost-effectively' = spot instances.

upvoted 1 times

amministrazione 10 months ago

D. Split the 12 instances across three Availability Zones in the chosen AWS Region. Run three instances in each Availability Zone as On-Demand Instances with Capacity Reservations. Run one instance in each Availability Zone as a Spot Instance.

upvoted 1 times

Helpnosense 1 year ago

Selected Answer: C

I vote C.
The 65% of scheduled jobs is the portion of the total work load. I don't believe it's SLA since SLA will be 99.99% or more. The jobs is hourly from 0.3 to 2 hours. There are 12 servers on prem. If the number of jobs per server can handle is N. Then to cover the worst situation that all the jobs run 2 hours, by given 12 servers and tight SLA, the number of hourly jobs is 12 / 2 = 6N. Answer C has 6 servers and since the number of job per server is N

then 6 server can handle 6N jobs match the hourly job number 6N.
2 ec2 with saving plan + 2 spot instances is more cost effective than 3 ec2 with capacity plan(not saving a penny by capacity reservation plan) + 1 spot instance.
upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: D

D is more cost-effective than C
upvoted 2 times

atirado 1 year, 6 months ago

Selected Answer: D

Option A - This option might not work: it might not provide sufficient processing capacity for the batch jobs to meet the SLAs during outages. Moreover, 4 servers will not provide sufficient capacity to meet the SLAs of batch jobs

Option B - This option might not work: In case of an outage affecting the On-Demand instances there might not be enough processing capacity to meet batch job SLAs
Option C - This option will not meet the requirement not to make any long-term commitments

Option D - This option will work: There is enough sufficient processing capacity to meet the SLAs of batch jobs and keep processing One-off jobs
upvoted 2 times

subbupro 1 year, 6 months ago

D would be perfect, because it requires more cpu usage, we should have more capacity CPU .
upvoted 1 times

edder 1 year, 7 months ago

Selected Answer: D

The answer is D.
Since it originally had a completely redundant configuration, it is thought that scheduled tasks are executed on 4 machines and user tasks are executed on 2 machines.
A,B: Requirements cannot be met when a specific region falls.
C: No Savings Plan required.
D: Even if a specific region goes down, 6 machines will be maintained, so service can be maintained.
upvoted 2 times

Russs99 1 year, 10 months ago

Selected Answer: D

About 65% or about 8 instances have to run at the same time to meet the SLA.
upvoted 3 times

ggrodskiy 1 year, 11 months ago

Correct C.
Option D is incorrect because running three instances in each Availability Zone as On-Demand Instances with Capacity Reservations will increase the cost of the solution without providing any additional benefit. Capacity Reservations are not necessary when using a Savings Plan, as they both offer guaranteed capacity at a discounted pricehttps://docs.aws.amazon.com/whitepapers/latest/how-aws-pricing-works/amazon-ec2.html. Also, running only one instance in each Availability Zone as a Spot Instance will not provide enough capacity for the user jobs that account for 35% of system usage.
upvoted 4 times

joefromnc 1 year, 10 months ago

Can't be C it says it can't require long term commitment. Savings plans like reserved instance require long term commitments with a contract.
upvoted 4 times

awsrd2023 1 year, 11 months ago

Selected Answer: D

D. 3 AZ (Redundancy), 3 EC2 in each AZ as on demand and 1 spot (addresses SLA in 65/35 ratio)
Ruling out Factors:
A. Only 2 AZ (low redundancy), all EC2 in capacity reservation (Not Cost effective as SLA requirement is in 65/35 ratio).
B. All 4 on-demand in 1 AZ (low redundancy), rest spot (Will efect tight SLA - is actually 35/65 instead of 65/35).
C. Savings Plan (Against no long term commitments requirement).
upvoted 3 times

NikkyDicky 2 years ago

Selected Answer: D

D

1 - need capacity reservation

2 - need to cover 65% with HA

upvoted 1 times

aca1 2 years, 1 month ago

Selected Answer: D

Just D is the right one. We need to garantee 65% (about 8 instances of 12) of capacity for the SLA, so 9 can do it and then let the others as spot. Another point Saving Plans need commitment "Savings Plans are a flexible pricing model that offer low prices on Amazon EC2, AWS Lambda, and AWS Fargate usage, in exchange for a commitment to a consistent amount of usage (measured in $/hour) for a 1 or 3 year term" - https://aws.amazon.com/savingsplans/compute-pricing/

upvoted 3 times

gameoflove 2 years, 1 month ago

Selected Answer: C

Voted C, the reason for this option is Spot Instance which is truely cost saving when we are performing Batch jobs and if you plan the cost properly this is best solution

upvoted 1 times

Maria2023 2 years, 2 months ago

Selected Answer: D

65% SLA can be reached only on answer D. Yeah - 9 instances are a bit too much but that's the only answer that meets the SLA

upvoted 1 times

Question #26 *Topic 1*

A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an encrypted file in Amazon S3. For the next version of the application, the security engineer wants to implement the following application design changes to improve security:
The database must use strong, randomly generated passwords stored in a secure AWS managed service.
The application resources must be deployed through AWS CloudFormation.
The application must rotate credentials for the database every 90 days.
A solutions architect will generate a CloudFormation template to deploy the application.
Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

A. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Specify a Secrets Manager RotationSchedule resource to rotate the database password every 90 days.

B. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Create an AWS Lambda function resource to rotate the database password. Specify a Parameter Store RotationSchedule resource to rotate the database password every 90 days.

C. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.

D. Generate the database password as a SecureString parameter type using AWS Systems Manager Parameter Store. Specify an AWS AppSync DataSource resource to automatically rotate the database password every 90 days.

Suggested Answer: *B*

Community vote distribution

A (100%)

Untamables Highly Voted 2 years, 6 months ago

Selected Answer: A

A

https://docs.aws.amazon.com/secretsmanager/latest/userguide/cloudformation.html

Option B is wrong. The ParameterStore::RotationSchedule resource does not exist in CloudFormation.

Option C is wrong. It does not meet the requirement because it does not use CloudFormation.

Option D is wrong. The AWS::AppSync::DataSource resource is what to create data sources for resolvers in AWS AppSync to connect to.

upvoted 18 times

OnePunchExam 2 years, 2 months ago

Agree with A but I want to nitpick on this reply "The ParameterStore::RotationSchedule resource does not exist in CloudFormation". It is technically more correct to say ParameterStore does not support automated rotation of secrets instead of saying ParameterStore::RotationSchedule is not supported by CF.

upvoted 9 times

karma4moksha Highly Voted 2 years, 1 month ago

Ans A but answer is badly phrased. Why is the Lambda needed ?

Refer docs: Some services offer managed rotation, where the service configures and manages rotation for you. With managed rotation, you don't use an AWS Lambda function to update the secret and the credentials in the database. The following services offer managed rotation:

Amazon RDS offers managed rotation for master user credentials. For more information, see Password management with Amazon RDS and AWS Secrets Manager in the Amazon RDS User Guide.

upvoted 14 times

soulation 4 months ago

Read it again more carefully: "offers managed rotation for master user credentials"

This is for application credential. Beside, even rotation for master has limitations:

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-secrets-manager.html

upvoted 1 times

ftaws 1 year, 5 months ago

I agree with you. Secret Manager support to rotate credentials.

upvoted 3 times

amministrazione Most Recent 10 months ago

A. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Specify a Secrets Manager RotationSchedule resource to rotate the database password every 90 days.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: A

Secrets Manager ($$$): Automatic rotation of secrets with AWS Lambda // SSM Parameter Store ($): No secret rotation (can enable rotation using Lambda triggered by EventBridge) --> more overhead even if it is cheaper ==> Answer A

upvoted 1 times

ivarnarik1 1 year, 1 month ago

Correct Answer: A

Cloudformation template::systems manager has no resource called: RotationSchedule. where as Cloudformation template::secrets manager Indeed has a resource called: RotationSchedule.

Therefore the correct answer is A only.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A is the correct answer

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: A

Option A is the most straightforward and provides the least amount of operational overhead because it leverages AWS Secrets Manager's native capabilities for secret rotation. This eliminates the need for custom rotation logic or external triggers for rotation, unlike the other options that either rely on AWS Systems Manager Parameter Store (which does not have built-in secret rotation capabilities like Secrets Manager) or require additional resources such as Amazon EventBridge or AWS AppSync for triggering rotations, which complicates the architecture and increases operational overhead.

Therefore, Option A is the correct choice as it directly addresses all the specified requirements using the intended features of AWS services, ensuring security and efficiency with minimal operational complexity.

upvoted 3 times

AimarLeo 1 year, 4 months ago

OK.. A ..but.. lambda to rotate for Secret Managers ? it does rotation natively ! why is that

upvoted 4 times

atirado 1 year, 6 months ago

Selected Answer: A

Option A - This option will work: This option takes advantage of the Automatic Rotation feature in Secrets Manager which reduces operational overhead during secret rotation, i.e. CloudTrail will show a secret was rotated

Option B - This option will not work: Parameter Store does not have a feature called RotationSchedule

Option C - This option might work but increases overhead: Rotation will be triggered on the 90 day schedule but more work will be necessary to validate the secret was rotated and tested, i.e. CloudTrail logs will only show a lambda function was triggered

Option D - This option will not work: Parameter Store does not have a feature called RotationSchedule

upvoted 4 times

shaaam80 1 year, 6 months ago

Selected Answer: A

Answer A. Password rotation -> Secrets Manager

upvoted 1 times

whenthan 1 year, 10 months ago

Selected Answer: A

Which resources specified in the CloudFormation template will meet the security engineer's requirements with the LEAST amount of operational overhead?

use https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-secretsmanager-rotationschedule.html

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

All - I feel the answer is A but why does it says Correct Answer "B" - What is the rationale behind B, can anyone explain. I am so confused??

upvoted 2 times

SuperDuperPooperScooper 1 year, 8 months ago

The answers shown as correct are almost never the right ones on these test dumps, just pay attention to what was most voted and the discussions in the comments

upvoted 4 times

chico2023 1 year, 10 months ago

Selected Answer: A

Answer: A

upvoted 1 times

NikkyDicky 2 years ago

Selected Answer: A

it's n A

upvoted 1 times

rtguru 2 years, 1 month ago

A poorly phrased but seems to be the best option in this scenario

upvoted 1 times

gameoflove 2 years, 1 month ago

Selected Answer: A

AWS Secret Manager is the best option for Password safety and option fulfill all the requirement

upvoted 1 times

chiplyti 2 years, 2 months ago

Selected Answer: A

A correct

upvoted 1 times

Question #27 *Topic 1*

A company is storing data in several Amazon DynamoDB tables. A solutions architect must use a serverless architecture to make the data accessible publicly through a simple API over HTTPS. The solution must scale automatically in response to demand.
Which solutions meet these requirements? (Choose two.)

A. Create an Amazon API Gateway REST API. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.

B. Create an Amazon API Gateway HTTP API. Configure this API with direct integrations to Dynamo DB by using API Gateway's AWS integration type.

C. Create an Amazon API Gateway HTTP API. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.

D. Create an accelerator in AWS Global Accelerator. Configure this accelerator with AWS Lambda@Edge function integrations that return data from the DynamoDB tables.

E. Create a Network Load Balancer. Configure listener rules to forward requests to the appropriate AWS Lambda functions.

Suggested Answer: *CD*

Community vote distribution

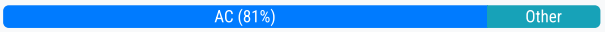

Untamables Highly Voted 2 years, 6 months ago

Selected Answer: AC

A and C.
API Gateway REST API can invoke DynamoDB directly.
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-overview-developer-experience.html
upvoted 29 times

ixdb 1 year, 6 months ago

CD is right.
While this option A works for private access, it does not support public access as DynamoDB tables are not publicly accessible by default.
upvoted 2 times

Impromptu 1 year, 6 months ago

Option A has the ability to specify an execution role. This IAM role should have the GetItem/PutItem permissions for the given DynamoDB table. That way you can have access to your private table via the DynamoDB API while your API Gateway is publicly available.
So I agree with A and C
upvoted 3 times

jpa8300 1 year, 6 months ago

You cannot choose A and C, you choose A OR C, one excludes the other. When a question says to choose two answers, one shall complement the other.
I agree that the API can integrate directly with DynamoDB, but if I have to choose two answers that complement each other, the A option cannot go with any of the others.
Saying that, the only possible choices should be C and D, you create the Lambda functions to integrate with Dynamodb and then deploy them at Edge, as extra to improve performance and latency you use Global Accelerator. Yes, it is true that this is not a requirement, but it is good to have.
upvoted 4 times

cegama543 8 months ago

A and C are compatibles. A for DynamoDB integration and C for escalation
upvoted 3 times

atirado Highly Voted 1 year, 6 months ago

Selected Answer: AC

Option A - This option might work: REST APIs can run over HTTPS and the integration type DynamoDB is possible

Option B - This option will not work: HTTP APIs do not support integration types for DynamoDB

Option C - This option will work: HTTP APIs support integration with Lambda functions

Option D - This option will not work: Lambda@Edge is a function of CloudFront

Option E - This option will not work: NLB Target groups can target Lambda functions however NLBs are not a Serverless solution (They are deployed on VPCs).

upvoted 13 times

CharChe Most Recent 1 week, 6 days ago

Selected Answer: CD

A should be ruled out because although REST API in API GW can be used to access DyanmaDB directly, accessing multiple DynamoDB tables by one single API needs a Lambda function to be created.

upvoted 1 times

TariqKipkemei 2 months, 2 weeks ago

Selected Answer: AC

REST API in API Gateway supports AWS service integrations such HTTP endpoints, Lambda functions, or other AWS services including direct calls to DynamoDB without Lambda.

HTTP APIs in API Gateway only support integrations to lambda functions.

Both options are fully serverless, support HTTPS, and scale automatically.

upvoted 1 times

Sin_Dan 8 months, 2 weeks ago

A. Amazon API Gateway can be configured as a REST API with direct integration to DynamoDB. This is done using API Gateway's AWS integration type, allowing direct interaction with DynamoDB without needing a Lambda function in between.

C. API Gateway HTTP API can be used to route requests to AWS Lambda functions. The Lambda functions can then interact with DynamoDB to retrieve or modify data and return it to the client through the API.

upvoted 1 times

amministrazione 10 months ago

A. Create an Amazon API Gateway REST API. Configure this API with direct integrations to DynamoDB by using API Gateway's AWS integration type.

C. Create an Amazon API Gateway HTTP API. Configure this API with integrations to AWS Lambda functions that return data from the DynamoDB tables.

upvoted 1 times

jyrajan69 10 months, 1 week ago

On the fact of simplicity it looks like BC, but with C there is an issue of Lambda fetching data, question does not indicate fetching, only put. So it looks like AB

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: AC

A and C

upvoted 1 times

Russs99 1 year, 3 months ago

Selected Answer: CD

The solutions that meet the requirements of using a serverless architecture to make the data accessible publicly through a simple API over HTTPS and scaling automatically in response to demand are: C AND D

upvoted 1 times

Russs99 1 year, 3 months ago

Actually, Option D is out, reason: you cannot use AWS Lambda@Edge with Global accelerator

upvoted 1 times

JOKERO 1 year, 3 months ago

a, c

https://medium.com/brlink/rest-api-just-with-apigateway-and-dynamodb-8a9b0cd76b7a

upvoted 1 times

anubha.agrahari 1 year, 3 months ago

Selected Answer: A

API Gateway REST API can invoke DynamoDB directly.

upvoted 1 times

DmitriKonnovNN 1 year, 4 months ago

Sometimes when multiple answers are required, they're supposed to complement each other, but sometimes these have to be just 2 valid but independent solutions... Well API GW with Rest endpoint is a valid solution, since it's had DynamoDB proxy integration lately. We use it in production, and it's a good fit, if you want to have a lot of control and features in your API GW and no lambda functions in between, reason being VTL supports a big set of mutations which is enough to us.
On the flip side, since we're forced to use a combination, then CD is the right answer.
In terms of simplicity, it is the question, what you consider simple. API GW REST endpoint is considered simple, because it provides caching, api keys, usage plans, rate limiting, authorization, deployment stages etc. out of the box. So the plethora of out-of-the-box features is rather simple than implementing them oneself.

upvoted 1 times

ninomfr64 1 year, 6 months ago

Selected Answer: BC

Not E as I think NLB listener rules don't provide the required capability to to forward requests to the appropriate Lambda (you need to have and ALB)
Not D as Lambda@Edge is a CloudFront feature

A, B and C they all works here however the question requires "a simple API over HTTPS". Both REST APIs and HTTP APIs are RESTful API products. REST APIs support more features than HTTP APIs, while HTTP APIs are designed with minimal features so that they can be offered at a lower price.
Thus I would go for B and C

upvoted 1 times

ninomfr64 1 year, 6 months ago

My answer is wrong, double check that DynamoDB is not supported as first-class integration with API Gateway as per doc
https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-develop-integrations-aws-services-reference.html

Thus the correct answer is A and C

upvoted 2 times

subbupro 1 year, 6 months ago

C and D is the correct option
1) C- Need server less architecture so need to use lamda function instead of REST API
2) D - Global accelerator work with lamda edge would be best the option compare to NLB for auto scale up and down. It has static address and fixed entry point if we deply multiple region.

upvoted 2 times

Hit1979 1 year, 7 months ago

Selected Answer: CE

REST API - is not simple and limitaiton around scability. NLB with listener rules can be used to forward request based on specified conditions to appropriate AWS lambda function

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: AC

lambda can have https endpoints available

upvoted 1 times

rodrod 1 year, 9 months ago

Selected Answer: BC

I've read similar questions previously, keyword is "simple API".
REST API adds more features than HTTP API and is consider "more" complex.
So it has to be HTTP just for that reason.

You can use API Gateway (HTTP)->dynamodb:
https://aws.amazon.com/fr/blogs/compute/using-amazon-api-gateway-as-a-proxy-for-dynamodb/

so B and C

upvoted 3 times

sonyaws 1 year, 7 months ago

BC
HTTP API support AWS Integrations + Simple
https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-vs-rest.html

upvoted 2 times

Question #28 *Topic 1*

A company has registered 10 new domain names. The company uses the domains for online marketing. The company needs a solution that will redirect online visitors to a specific URL for each domain. All domains and target URLs are defined in a JSON document. All DNS records are managed by Amazon Route 53.
A solutions architect must implement a redirect service that accepts HTTP and HTTPS requests.
Which combination of steps should the solutions architect take to meet these requirements with the LEAST amount of operational effort? (Choose three.)

A. Create a dynamic webpage that runs on an Amazon EC2 instance. Configure the webpage to use the JSON document in combination with the event message to look up and respond with a redirect URL.

B. Create an Application Load Balancer that includes HTTP and HTTPS listeners.

C. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.

D. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.

E. Create an Amazon CloudFront distribution. Deploy a Lambda@Edge function.

F. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.

Suggested Answer: *BCF*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: CEF

C: By creating an AWS Lambda function, the solution architect can use the JSON document to look up the target URLs for each domain and respond with the appropriate redirect URL. This way, the solution does not need to rely on a web server to handle the redirects, which reduces operational effort.

E: By creating an Amazon CloudFront distribution, the solution architect can deploy a Lambda@Edge function that can look up the target URLs for each domain and respond with the appropriate redirect URL. This way, CloudFront can handle the redirection, which reduces operational effort.

F: By creating an SSL certificate with ACM and including the domains as Subject Alternative Names, the solution architect can ensure that the redirect service can handle both HTTP and HTTPS requests, which is required by the company.

upvoted 37 times

Shahul75 2 years, 4 months ago

SAN cannot handle redirects. We need to do http - https

upvoted 1 times

masetromain 2 years, 5 months ago

A and B are not the right answer because they would require configuring and maintaining a web server to handle the redirects, which would increase operational effort.
D is not the right answer because it would require creating an API Gateway API, which increases operational effort.

upvoted 7 times

Arnaud92 2 years, 3 months ago

Wrong for B, Lambda can be an ALB target, you do not need web server

upvoted 9 times

chathur Highly Voted 2 years, 1 month ago

Selected Answer: BCF

If you go with a Cloudfront what is the origin? Lambda@edge is not origin. The function mentioned in C is Lambda and in E it says about Lambda@edge, which are two things. If you handle redirect from the Lambda@edge in CF there is no need of the Lambda you wrote in Answer C.

MY Answer:
Create an ALB with HTTP and HTTPS listeners (B), Use the TLS cert created in F for the HTTPS listener. As the backend for the ALB write a Lambda with endpoint mapping JSON (C)

Is this full serverless? No, but you do not have to worry about scaling or operational overhead, AWS Handles everything for us.
upvoted 36 times

dubyaF 1 year, 6 months ago
This is the only answer that is completed by using all three options selected BCF. F is mandatory to resolve the marketing domains URLs that are HTTPS. So B and C then work together to redirect to those URLs as a full solution like https://aws.amazon.com/ko/blogs/networking-and-content-delivery/automating-http-s-redirects-and-certificate-management-at-scale/

E may have partial potential to do something, but you have no origin with it - and what would the origin be?
With BCF you hit the ALB get a redirect as a result of the Marketing URL and your done-- its a complete redirect solution which is what the whole requirement is.
upvoted 6 times

ex_ample Most Recent 1 week ago
Selected Answer: CEF
CEF
C: Read json and redirect
E: Handle HTTP/HTTPS (intercept request)
F: Required for multiple Certificate management (HTTPS)

WRONG
A: EC2 too much
B: ALB handles HTTP/HTTPS BUT no dynamic lookup using JSON
D: More operation effort comapre to CloudFront
upvoted 1 times

Kaps443 3 weeks, 2 days ago
Selected Answer: CDF
Simple, low-maintenance, HTTPS support, no compute needed
upvoted 1 times

CAIYasia 2 months ago
Selected Answer: CDF
Request → Route 53DNS resolve → API GatewayDomain name + ACM cert. → Lambdaredirect
upvoted 2 times

ed605fe 2 months, 1 week ago
Selected Answer: BCF
Best match
upvoted 1 times

codeScalable 3 months, 2 weeks ago
Selected Answer: BCF
BCF. You already have a lambda function that does the processing. You don't need Lambda@Edge anymore
upvoted 1 times

3a05e15 4 months, 2 weeks ago
Selected Answer: CDF
Option E involves CloudFront with Lambda@Edge, which is more complex to manage than API Gateway for this use case.
upvoted 2 times

altonh 5 months, 3 weeks ago
Selected Answer: BCF
CEF - an overkill
CDF - D says that you should create a custom domain. You need to perform several steps just for this custom domain so it is a more complicated setup.
upvoted 1 times

altonh 5 months, 3 weeks ago
Also, API Gateway only exposes HTTPS endpoints.
upvoted 1 times

thanhpolimi 6 months ago

Selected Answer: CDF

CDF is the best combination.

upvoted 1 times

pk0619 6 months, 2 weeks ago

Selected Answer: CDF

Lambda@Edge can perform redirection at CloudFront's edge locations, but this is more complex and costly for a simple redirect service.
CloudFront is overkill for this use case as it primarily serves content delivery purposes.

upvoted 1 times

wem 6 months, 3 weeks ago

Selected Answer: CDF

C. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
Why:
AWS Lambda provides a serverless solution to handle the redirect logic. It can read the domain and use the JSON document to determine the appropriate target URL.
This reduces the operational overhead of managing infrastructure like EC2 instances.

D. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
Why:
API Gateway provides a fully managed way to route HTTP(S) requests to the Lambda function.It supports custom domains, allowing you to set up a single entry point for the redirect service, with SSL/TLS support.

F. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.
Why:
ACM allows you to issue free SSL/TLS certificates for all 10 domains.
This ensures HTTPS support for the redirect service without manual certificate management, aligning with the requirement to handle both HTTP and HTTPS.

upvoted 4 times

karlgrosz 6 months, 4 weeks ago

Selected Answer: CDF

E: Lambda@Edge is more suitable for content delivery network (CDN)-focused use cases, not for redirect services. It also requires additional setup and complexity compared to the API Gateway approach.

upvoted 1 times

manngupta007 7 months ago

Selected Answer: CDE

Deploy Lambda function that looks up and returns the appropriate URL from the JSON document.
Publish Lambda using API Gateway with a custom domain name for your 10 registered domains.
Use CloudFront with Lambda@Edge to ensure fast and scalable global redirection.
This approach is serverless, scalable, and minimizes operational overhead.

upvoted 2 times

sergza 7 months ago

Selected Answer: CEF

I think it is CEF according to this https://aws.amazon.com/blogs/networking-and-content-delivery/handling-redirectsedge-part1/

upvoted 1 times

Tiger4Code 7 months ago

Selected Answer: CDF

The correct steps to implement the redirect service with the LEAST operational effort are:

C. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL.
Lambda allows serverless, scalable handling of requests, and is ideal for processing requests based on a JSON document.

D. Use an Amazon API Gateway API with a custom domain to publish an AWS Lambda function.
API Gateway provides a scalable interface to expose the Lambda function via HTTP/HTTPS and supports custom domains.

F. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names.
ACM simplifies the management of SSL certificates, ensuring HTTPS support for the custom domain.

upvoted 1 times

henrikhmkhitaryan59 7 months, 1 week ago

Selected Answer: BCF

B. Create an Application Load Balancer that includes HTTP and HTTPS listeners. (for secure and easy traffic management)

C. Create an AWS Lambda function that uses the JSON document in combination with the event message to look up and respond with a redirect URL. (for low-cost, serverless compute)

F. Create an SSL certificate by using AWS Certificate Manager (ACM). Include the domains as Subject Alternative Names. (for easy SSL management)

upvoted 1 times

Question #29 *Topic 1*

A company that has multiple AWS accounts is using AWS Organizations. The company's AWS accounts host VPCs, Amazon EC2 instances, and containers.
The company's compliance team has deployed a security tool in each VPC where the company has deployments. The security tools run on EC2 instances and send information to the AWS account that is dedicated for the compliance team. The company has tagged all the compliance-related resources with a key of "costCenter" and a value or "compliance".
The company wants to identify the cost of the security tools that are running on the EC2 instances so that the company can charge the compliance team's AWS account. The cost calculation must be as accurate as possible.
What should a solutions architect do to meet these requirements?

A. In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.

B. In the member accounts of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Schedule a monthly AWS Lambda function to retrieve the reports and calculate the total cost for the costCenter tagged resources.

C. In the member accounts of the organization activate the costCenter user-defined tag. From the management account, schedule a monthly AWS Cost and Usage Report. Use the tag breakdown in the report to calculate the total cost for the costCenter tagged resources.

D. Create a custom report in the organization view in AWS Trusted Advisor. Configure the report to generate a monthly billing summary for the costCenter tagged resources in the compliance team's AWS account.

Suggested Answer: *A*

Community vote distribution

A (95%) 5%

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

Answer A : because we do not depend on the users, I prefer management account

Option C or A would be the correct answer. In option C, the solution architect would activate the costCenter user-defined tag in the member accounts of the organization, and then schedule a monthly AWS Cost and Usage Report from the management account to retrieve the reports and calculate the total cost for the costCenter tagged resources. In option A, the management account of the organization would activate the costCenter user-defined tag and configure monthly AWS Cost and Usage Reports to be saved to an Amazon S3 bucket in the management account. Then, use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources. Both options would allow the company to accurately identify the cost of the security tools running on the EC2 instances and charge the compliance team's AWS account.

upvoted 20 times

dkx 2 years ago

Only a management account in an organization and single accounts that aren't members of an organization have access to the cost allocation tags manager in the Billing and Cost Management console.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html

upvoted 13 times

chathur 2 years, 1 month ago

User-defined tags can not be allowed from management accounts in AWS Organization. It must done from the management Account.

upvoted 2 times

Reval 12 months ago

Did you mean from member account? in this sentence "User-defined tags can not be allowed from management accounts in AWS Organization."

upvoted 1 times

Untamables Highly Voted 2 years, 6 months ago

Selected Answer: A

I vote A.
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/configurecostallocreport.html
upvoted 6 times

jimee11 Most Recent 1 month, 3 weeks ago

Selected Answer: A

Cost tags are activated in the Management console.
upvoted 1 times

Tiger4Code 7 months ago

Selected Answer: A

A. In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.
upvoted 1 times

amministrazione 10 months ago

A. In the management account of the organization, activate the costCenter user-defined tag. Configure monthly AWS Cost and Usage Reports to save to an Amazon S3 bucket in the management account. Use the tag breakdown in the report to obtain the total cost for the costCenter tagged resources.
upvoted 1 times

Jason666888 11 months ago

Selected Answer: A

The most ideal way to get this job done is to use: AWS Cost Explorer

But among all the given options, we should go with option A, as the user defined tag can only be managed in management account
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A is correct
upvoted 1 times

subbupro 1 year, 6 months ago

A is ccorect, we need to login to management account to create
upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: A

yes, you need to activate cost allocation tags before using, you can do this the same place where you would like to see your reports - management account
upvoted 2 times

whenthan 1 year, 8 months ago

Selected Answer: C

lines up correctly
activate tag in member accounts and generating AWS CUR from management account (has ability to see costs across all member accounts) and Tag breakfdown in report
upvoted 1 times

imvb88 1 year, 9 months ago

Selected Answer: A

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/activating-tags.html
"For tags to appear on your billing reports, you must activate them."

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.html
"Only a management account in an organization and single accounts that aren't members of an organization have access to the cost allocation tags manager in the Billing and Cost Management console."
-> eliminate B,C. D is not relevant
upvoted 2 times

whenthan 1 year, 9 months ago

Selected Answer: A

https://docs.aws.amazon.com/whitepapers/latest/tagging-best-practices/building-a-cost-allocation-strategy.html

upvoted 1 times

bur4an 1 year, 10 months ago

Selected Answer: A

Only a management account in an organization and single accounts that aren't members of an organization have access to the cost allocation tags manager in the Billing and Cost Management console.

upvoted 3 times

NikkyDicky 2 years ago

it's an A

upvoted 1 times

rtguru 2 years, 1 month ago

I go with D

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: A

Cost center tag int he management account.

upvoted 1 times

kiran15789 2 years, 3 months ago

Selected Answer: A

Management account for reports

upvoted 1 times

Question #30 *Topic 1*

A company has 50 AWS accounts that are members of an organization in AWS Organizations. Each account contains multiple VPCs. The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account. Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (Choose two.)

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager.

B. From the management account, share the transit gateway with member accounts by using an AWS Organizations SCP.

C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.

D. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a peering transit gateway attachment in a member account. Share the attachment with the transit gateway in the management account by using a transit gateway service-linked role.

E. From the management account, share the transit gateway with member accounts by using AWS Service Catalog.

Suggested Answer: *AC*

Community vote distribution

AC (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: AC

Option A is sharing the transit gateway with member accounts by using AWS Resource Access Manager, which allows the management account to share resources with member accounts. Option C is launching an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account, and associates the attachment with the transit gateway in the management account by using the transit gateway ID. This automation of creating a new VPC and transit gateway attachment in new member accounts can help to streamline the process and reduce operational effort.

upvoted 23 times

jainparag1 1 year, 7 months ago

Precisely!

upvoted 1 times

ausl Most Recent 1 month, 3 weeks ago

Selected Answer: AC

From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager

upvoted 1 times

Tiger4Code 7 months ago

Selected Answer: AC

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager. Most Voted
C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.
You want to associate the gateway attachment with the transit gateway that you already shared using RAM

upvoted 1 times

amministrazione 10 months ago

A. From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager.
C. Launch an AWS CloudFormation stack set from the management account that automatically creates a new VPC and a VPC transit gateway attachment in a member account. Associate the attachment with the transit gateway in the management account by using the transit gateway ID.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: AC

AC are correct

upvoted 1 times

[Removed] 1 year, 6 months ago

Selected Answer: AC

I am working on a project doing the exact same thing :D

upvoted 2 times

rlf 1 year, 8 months ago

AC.

https://aws.amazon.com/ko/blogs/networking-and-content-delivery/automating-aws-transit-gateway-attachments-to-a-transit-gateway-in-a-central-account/

https://cloudjourney.medium.com/aws-ram-and-transit-gateway-8ac230f298e8

upvoted 1 times

Simon523 1 year, 9 months ago

Selected Answer: AC

You can use AWS Resource Access Manager (RAM) to share a transit gateway for VPC attachments across accounts or across your organization in AWS Organizations.

upvoted 1 times

NikkyDicky 2 years ago

AC of course

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: AC

AC are my choice.

upvoted 2 times

zozza2023 2 years, 5 months ago

Selected Answer: AC

A and C are the answer for me

upvoted 2 times

Untamables 2 years, 6 months ago

Selected Answer: AC

A & C

https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.html

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachment.html

upvoted 2 times

masetromain 2 years, 6 months ago

Selected Answer: AC

https://www.examtopics.com/discussions/amazon/view/60090-exam-aws-certified-solutions-architect-professional-topic-1/

upvoted 3 times

Question #31 *Topic 1*

An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace. The company uses an AWS Organizations account structure with full features enabled, and has a shared services account in each organizational unit (OU) that will be used by procurement managers. The procurement team's policy indicates that developers should be able to obtain third-party software from an approved list only and use Private Marketplace in AWS Marketplace to achieve this requirement. The procurement team wants administration of Private Marketplace to be restricted to a role named procurement-manager-role, which could be assumed by procurement managers. Other IAM users, groups, roles, and account administrators in the company should be denied Private Marketplace administrative access.
What is the MOST efficient way to design an architecture to meet these requirements?

A. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the PowerUserAccess managed policy to the role. Apply an inline policy to all IAM users and roles in every AWS account to deny permissions on the AWSPrivateMarketplaceAdminFullAccess managed policy.

B. Create an IAM role named procurement-manager-role in all AWS accounts in the organization. Add the AdministratorAccess managed policy to the role. Define a permissions boundary with the AWSPrivateMarketplaceAdminFullAccess managed policy and attach it to all the developer roles.

C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.

D. Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an SCP in Organizations to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Apply the SCP to all the shared services accounts in the organization.

Suggested Answer: *D*

Community vote distribution

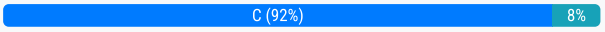

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The most efficient way to design an architecture to meet these requirements is option C. By creating an IAM role named procurement-manager-role in all the shared services accounts in the organization and adding the AWSPrivateMarketplaceAdminFullAccess managed policy to the role, the procurement managers will have the necessary permissions to administer Private Marketplace. Then, by creating an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role and another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization, the company can restrict access to Private Marketplace administrative access to only the procurement managers.

upvoted 16 times

SK_Tyagi 1 year, 10 months ago

The catch is the "Create an organization root-level SCP to deny permissions". I'd refrain from creating a root-level SCP

upvoted 3 times

amministrazione Most Recent 10 months ago

C. Create an IAM role named procurement-manager-role in all the shared services accounts in the organization. Add the AWSPrivateMarketplaceAdminFullAccess managed policy to the role. Create an organization root-level SCP to deny permissions to administer Private Marketplace to everyone except the role named procurement-manager-role. Create another organization root-level SCP to deny permissions to create an IAM role named procurement-manager-role to everyone in the organization.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: C

Not D, why ? : D. Placing the procurement-manager-role in developer accounts with full Private Marketplace admin access increases the risk of mismanagement. Additionally, applying an SCP only to shared services accounts does not adequately restrict access across the entire organization.

upvoted 1 times

cnethers 1 year ago

Why C is right and D is wrong....
Focus on the end of the question :
Other IAM users, groups, roles, and account administrators in the company should be denied Private Marketplace administrative access.
What is the MOST efficient way to design an architecture to meet these requirements?

Who should be excluded? Other IAM users, groups, roles, and account administrators in the company
What is the MOST efficient way? Apply SCP at the root level
D is more work than C, this is a good reason to choose C over D
upvoted 1 times

Chakanetsa 1 year, 1 month ago

Selected Answer: C

C. Most efficient and secure:

Creating the procurement-manager-role in shared services accounts limits its scope to specific OUs, aligning with the organizational structure.
Granting AWSPrivateMarketplaceAdminFullAccess to this role provides the necessary permissions for managing Private Marketplace within the OU.
An organization root-level SCP denying Private Marketplace administration to everyone except the procurement-manager-role ensures centralized control and restricts unauthorized access.
Another SCP preventing the creation of the procurement-manager-role outside of shared services accounts adds an extra layer of security.
upvoted 1 times

anubha.agrahari 1 year, 3 months ago

Selected Answer: C

C, D doesn't make sense.
upvoted 1 times

ninomfr64 1 year, 6 months ago

Selected Answer: C

Not A as it does not implement the requirement to enforce procurement managers to use the shared services account in each organizational unit
Not B as this would allow developers to administer private market place
not D as this would allow developers to administer private market place

C is correct as it configure the required role (with required permission) only in the shared service account, uses an SCP to deny private market place management to to everyone except the role named procurement-manager-role and another SCP to prevent creating a role nmaed procurement-manager-role
upvoted 2 times

ninomfr64 1 year, 6 months ago

Actually D would to the job, but creating a role in every account is nt strictly necessary and would cause more work
upvoted 1 times

subbupro 1 year, 6 months ago

C is the better one than D . because we need to apply scp to the root level with deny policy is the best practices. create the role and apply to each account is not a correct way and it is overhead to the adminstrator.
upvoted 2 times

severlight 1 year, 7 months ago

Selected Answer: C

look on whenthan's answer
upvoted 1 times

whenthan 1 year, 8 months ago

Selected Answer: C

creation of role in all shared services
adding required policy to the role
creation of org root-level to guardrail who can have those privileges
creation of SCP to close out workaround of creation of another role with same access
upvoted 3 times

Tarun4b7 1 year, 9 months ago

Selected Answer: D

C and D options are most relevant. Once you create a role, you cannot create another role with same name. So option C doesn't make sense. So my answer Option D

upvoted 2 times

Jassybanga 1 year, 4 months ago

i am on same page

upvoted 1 times

Jassybanga 1 year, 4 months ago

its C - the role should be in shared service accounts and not all accounts

upvoted 1 times

qxy 1 year, 9 months ago

Selected Answer: C

Clearly, it's C.

upvoted 1 times

Karamen 1 year, 10 months ago

Selected answer: C

option D: "Create an IAM role named procurement-manager-role in all AWS accounts that will be used by developers", the procurement-manager-role is used by manager not used by developers

upvoted 2 times

alicewsm 1 year, 8 months ago

the first sentense "An enterprise company wants to allow its developers to purchase third-party software through AWS Marketplace."

upvoted 1 times

jainparag1 1 year, 7 months ago

Developers has to ask procurement manager and not purchase by themselves.

upvoted 2 times

SorenBendixen 1 year, 10 months ago

Selected Answer: D

Its D - According to this : https://aws.amazon.com/blogs/awsmarketplace/controlling-access-to-a-well-architected-private-marketplace-using-iam-and-aws-organizations/

upvoted 2 times

SorenBendixen 1 year, 10 months ago

Its C. D is wrong - missed : "procurement-manager-role in all AWS accounts that will be used by DEVELOPERS"

upvoted 2 times

NikkyDicky 1 year, 11 months ago

Selected Answer: C

Its a C

upvoted 1 times

gd1 2 years ago

Selected Answer: C

C is correct-

upvoted 1 times

Maria2023 2 years ago

Selected Answer: C

D is a distractor since the developers do not need to administer the private marketplace. Plus that the procurement team acts only in the shared accounts. That leaves C as the only option

upvoted 4 times

Question #32 *Topic 1*

A company is in the process of implementing AWS Organizations to constrain its developers to use only Amazon EC2, Amazon S3, and Amazon DynamoDB. The developers account resides in a dedicated organizational unit (OU). The solutions architect has implemented the following SCP on the developers account:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowEC2",
            "Effect": "Allow",
            "Action": "ec2:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowDynamoDB",
            "Effect": "Allow",
            "Action": "dynamodb:*",
            "Resource": "*"
        },
        {
            "Sid": "AllowS3",
            "Effect": "Allow",
            "Action": "s3:*",
            "Resource": "*"
        }
    ]
}
```

When this policy is deployed, IAM users in the developers account are still able to use AWS services that are not listed in the policy.
What should the solutions architect do to eliminate the developers' ability to use services outside the scope of this policy?

A. Create an explicit deny statement for each AWS service that should be constrained.

B. Remove the FullAWSAccess SCP from the developers account's OU.

C. Modify the FullAWSAccess SCP to explicitly deny all services.

D. Add an explicit deny statement using a wildcard to the end of the SCP.

Suggested Answer: *A*

Community vote distribution

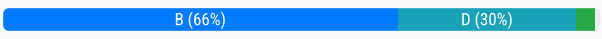

zhangyu20000 Highly Voted 2 years, 6 months ago

B is correct because default FullAWSAccess SCP is applied

upvoted 19 times

Six_Fingered_Jose Highly Voted 1 year, 9 months ago

Selected Answer: B

If you go to AWS management console and look up how SCP works, you will find that by default FullAWSAccess policy is attached to all OUs by default if you have SCP enabled.

upvoted 12 times

jainparag1 1 year, 7 months ago

That's correct. You can disable AWSFullAccess SCP from member accounts as long as you are replacing it with another policy with specific permissions required.

upvoted 4 times

jimee11 Most Recent 1 month, 3 weeks ago

Selected Answer: B

AWS advises you to replace the full access assigned by default in the SCP with your controlled SCP. Enabling SCPs gives you FullAWSAccess of the bat.

upvoted 1 times

diazed 2 months, 2 weeks ago

Selected Answer: B

B is correct. AWS Organizations attaches an AWS managed SCP named FullAWSAccess to every root, OU and account when it's created. This policy allows all services and actions. You can replace FullAWSAccess with a policy allowing only a set of services so that new AWS services are not allowed unless they are explicitly allowed by updating SCPs. For example, if your organization wants to only allow the use of a subset of services in your environment, you can use an Allow statement to only allow specific services. A policy combining the two statements might look like the following example, which prevents member accounts from leaving the organization and allows use of desired AWS services. The organization administrator can detach the FullAWSAccess policy and attach this one instead.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html

upvoted 2 times

GabrielShiao 3 months, 1 week ago

Selected Answer: A

I have to choose A although A is impractical. While most vote B, it is actually impossible since removing FullAWSAcess SCP from OU will deny all the services on the ous and accounts under the OU. The correct action is to remove FullAWSAccess SCP from the developer account.

upvoted 1 times

GabrielShiao 3 months, 1 week ago

Selected Answer: A

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html
If you removed the default SCP from the OU, you will be denied for these permission even you allowed in SCP on the account in OU.

upvoted 1 times

GabrielShiao 3 months, 1 week ago

Selected Answer: A

If you removed FullAWSAccess from developer accounts, I vote B, however, B is removing it from OU. Keep in mind every level of organization hierarchy must reside at least one SCP.

upvoted 1 times

konieczny69 8 months, 1 week ago

It can be as well handled with a or d, like
{
"Effect": "Deny",
"NotAction": [
"ec2:*",
"s3:*",
"dynamodb:*"
],
"Resource": "*"
}

upvoted 2 times

amministrazione 10 months ago

B. Remove the FullAWSAccess SCP from the developers account's OU.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: B

B. Remove the FullAWSAccess SCP from the developers account's OU.

Explanation:
FullAWSAccess SCP: By default, AWS Organizations attaches a FullAWSAccess SCP to all OUs and accounts, allowing access to all AWS services unless restricted by another SCP. If this SCP is still attached to the developers' OU, it will allow access to all services, regardless of the more restrictive SCP you have applied.

SCP Behavior: SCPs are evaluated in an "implicit deny" model. If an action is not explicitly allowed by the SCPs, it is implicitly denied. However, if multiple SCPs are attached and one allows an action (like FullAWSAccess), that action is permitted unless explicitly denied in another SCP.
upvoted 2 times

felon124 10 months, 3 weeks ago

Selected Answer: B

AWS Organizations attaches an AWS managed SCP named FullAWSAccess to every root, OU and account when it's created. This policy allows all services and actions. You can replace FullAWSAccess with a policy allowing only a set of services so that new AWS services are not allowed unless they are explicitly allowed by updating SCPs.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html
upvoted 1 times

8693a49 11 months ago

Selected Answer: D

Best practice would be to create an explicit deny statement. The reason is that other SCPs could be in effect, aside from AWSFullAccess, that could grant access to other services. If the goal is to deny access to any other service, then this must be made explicit.
upvoted 1 times

vip2 11 months ago

Selected Answer: B

B is correct
Remove from develop account OU --> implicitly deny all service -->add explicity 'allow' to restirct only allow related services in SCP.
upvoted 1 times

Moghite 11 months, 1 week ago

Selected Answer: D

{
"Sid": "ExplicitDeny",
"Effect": "Deny",
"NotAction": [
"ec2:*",
"dynamodb:*",
"s3:*"
],
"Resource": "*"
}
upvoted 2 times

Helpnosense 1 year ago

Selected Answer: D

FullAWSAccess SCP is inherited from root. Can't be removed from OU.
D is correct answer.
upvoted 2 times

sam2ng 10 months, 3 weeks ago

It can be, read "How SCPs work with Allow" in here it shows example:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html
upvoted 1 times

qaz12wsx 1 year, 2 months ago

Selected Answer: D

{
"Version": "2012-10-17",
"Statement": [
{
"Sid": "AllowEC2",
"Effect": "Allow",
"Action": "ec2:*",
"Resource": "*"
},

{
"Sid": "AllowDynamoDB",
"Effect": "Allow",
"Action": "dynamodb:*",
"Resource": "*"
},
{
"Sid": "AllowS3",
"Effect": "Allow",
"Action": "s3:*",
"Resource": "*"
},
{
"Sid": "ExplicitDeny",
"Effect": "Deny",
"NotAction": [
"ec2:*",
"dynamodb:*",
"s3:*"
],
"Resource": "*"
}
]
}

upvoted 4 times

Dgix 1 year, 3 months ago

Selected Answer: D

D - the alternative doesn't mention an ASG which must be taken as implied.

The other solutions are simply absurd:

A: The operational overhead is ENORMOUS. To those who think that "operational overhead" is only day-to-day maintenance: it is not. It encompasses ALL CHANGES to the infrastructure.

B: Kubernetes is the very definition of operational overhead. Always avoid unless there is an absolutely compelling reason to use it.

C: And what do you people think the function of the Lambda is? None.

D: This works and is the most straightforward as soon as you realise that the ASG is implied.

In the final analysis, this is another example of how AWS exam questions leave out information in order to trip you up.

upvoted 2 times

Question #33 *Topic 1*

A company is hosting a monolithic REST-based API for a mobile app on five Amazon EC2 instances in public subnets of a VPC. Mobile clients connect to the API by using a domain name that is hosted on Amazon Route 53. The company has created a Route 53 multivalue answer routing policy with the IP addresses of all the EC2 instances. Recently, the app has been overwhelmed by large and sudden increases to traffic. The app has not been able to keep up with the traffic.
A solutions architect needs to implement a solution so that the app can handle the new and varying load.
Which solution will meet these requirements with the LEAST operational overhead?

A. Separate the API into individual AWS Lambda functions. Configure an Amazon API Gateway REST API with Lambda integration for the backend. Update the Route 53 record to point to the API Gateway API.

B. Containerize the API logic. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Run the containers in the cluster by using Amazon EC2. Create a Kubernetes ingress. Update the Route 53 record to point to the Kubernetes ingress.

C. Create an Auto Scaling group. Place all the EC2 instances in the Auto Scaling group. Configure the Auto Scaling group to perform scaling actions that are based on CPU utilization. Create an AWS Lambda function that reacts to Auto Scaling group changes and updates the Route 53 record.

D. Create an Application Load Balancer (ALB) in front of the API. Move the EC2 instances to private subnets in the VPC. Add the EC2 instances as targets for the ALB. Update the Route 53 record to point to the ALB.

Suggested Answer: *D*

Community vote distribution

EricZhang Highly Voted 2 years, 6 months ago

Selected Answer: A

Serverless requires least operational effort.

upvoted 37 times

dqwsmwwvtgxwkvgcvc 1 year, 10 months ago

I guess multivalue answer routing in Route53 is not proper load balancing so replacing multivalue answer routing with ALB would proper balance the load (with minimal effort)

upvoted 4 times

lkyixoayffasdrlaqd 2 years, 4 months ago

How can this be the answer ?? It says: Separate the API into individual AWS Lambda functions. Can you calculate the operational overhead to do that?

upvoted 21 times

scuzzy2010 2 years, 2 months ago

Separating would be development overhead, but once done, the operational overheard (operational = ongoing day-to-day) will be the least.

upvoted 13 times

24Gel 1 year, 3 months ago

disagree, ASG in Option D, after set up, operational is not overheat as well

upvoted 1 times

24Gel 1 year, 3 months ago

i mean Option C not D

upvoted 1 times

24Gel 1 year, 3 months ago

never mind, A is simpler than C

upvoted 2 times

Jay_2pt0_1 2 years, 1 month ago

From any type of real-world perspective, this just can't be the answer IMHO. Surely AWS takes "real world" into account.

upvoted 1 times

jooncco Highly Voted 2 years, 5 months ago

Selected Answer: C

Suppose there are a 100 REST APIs (Since this application is monolithic, it's quite common).
Are you still going to copy and paste all those API codes into lambda?
What if business logic changes?
This is not MINIMAL. I would go with C.
upvoted 33 times

altonh 5 months, 3 weeks ago
Option C means your R53 is playing catch-up with your ASG. What happens if you scale down? Your clients will still have the terminated EC2 in their cache until the next TTL.
upvoted 1 times

chathur 2 years, 1 month ago
"Create an AWS Lambda function that reacts to Auto Scaling group changes and updates the Route 53 record. "
This does not make any sense, why do you need to change R53 records using a Lambda?
upvoted 1 times

Vesla 1 year, 10 months ago
Because if you have 4 ec2 in your ASG you need to have 4 records in domain name if ASG scale up to 6 for example you need 2 add 2 records more in domain name
upvoted 4 times

liquen14 1 year, 3 months ago
Too contrived in my opinion, and what about DNS caches in the clients?. You coul get stuck for a while with the previous list of servers. I think it's has to be A (but it would involve a considerable development effort) or D which is extremely easy to implement but and the same time it sounds a little bit fishy because they don't mention anything about ASG or scaling

I hate this kind of questions and I don't understand what kind of useful insight they provide unless they want us to become masters of the art of dealing with ambiguity
upvoted 3 times

cnethers 1 year ago
Agree that D does not scale to meet demand, it's just a better way to load balance which was being done at R53 before so the scaling issue has not been resolved.
Also agree A requires more dev effort and less ops effort, so I would have to lean to A...
Answer selection is poor IMO
upvoted 1 times

scuzzy2010 2 years, 4 months ago
It says "a monolithic REST-based API " - hence only 1 API. Initially I thought C, but I'll go with A as it says least operation overhead (not least implementation effort). Lambda has virtually no operation overhead compared to EC2.
upvoted 8 times

aviathor 1 year, 11 months ago
Answer A says "Separate the API into individual AWS Lambda functions." Makes me think there may be many APIs.

However, we are looking to minimize operational effort, not development effort...
upvoted 1 times

Jay_2pt0_1 2 years, 2 months ago
A monolithic REST api likely has a gazillion individual APIs. This refactor would not be a small one.
upvoted 5 times

jainparag1 1 year, 7 months ago
Dealing with business logic change is applicable to existing solution or any solution based on the complexity. Rather it's easier to deal when these are microservices. You shouldn't hesitate to refactor your application by putting one time effort (dev overhead) to save significant operational overhead on daily basis. AWS is pushing for serverless only for this.
upvoted 1 times

12db8b7 Most Recent 4 days, 11 hours ago
Selected Answer: D
I believe the issues the company is experiencing are directly related to the use of the Route 53 Multi-Value Answer routing policy without a proper load balancer. This approach relies on DNS to distribute traffic, which can result in uneven load distribution across EC2 instances. Since clients may randomly select the first IP returned by DNS, some instances might get overwhelmed while others remain underutilized.
In contrast, using an Application Load Balancer (ALB) alone but recomended with an Auto Scaling Group provides real-time traffic distribution,

automatic health checks, and the ability to scale out based on demand. All with minimal operational overhead. For this use case, where the application is monolithic and the main concern is handling unpredictable traffic surges, Option D (ALB) is the most reliable and scalable solution.
upvoted 1 times

sergza888 1 week, 6 days ago

Selected Answer: D

Lambda's refactorings, especially Labmdas that update DNS or EKS Ingress "Envoy" provisioning add to operational complexities, D it is quite simple and easy
upvoted 1 times

12db8b7 4 days, 11 hours ago

But u still don't upgrade the ability to scale with only a ALB, so I would go with C
upvoted 1 times

Monsterpuss 3 weeks ago

Selected Answer: C

My preference would be to go for option D as an ALB is a more elegant solution, but without an ASG behind it, there would still be problems. C is less elegant, but has the advantage of an ASG, even if the R53 update mechanism is messy.
upvoted 1 times

Kaps443 3 weeks ago

Selected Answer: D

ALB + Private EC2 + Route 53 → ALB
Perfect for immediate scaling needs, security, and minimal disruption.

A is incorrect
Requires completely re-architecting the monolithic API into functions — high development effort
Not the least operational overhead initially (only pays off long term)
upvoted 1 times

senlogan 2 months, 1 week ago

Selected Answer: A

Right answer would be ALB+Autoscaling. Go with A because the question is asking least operation overheard not least effort.
upvoted 2 times

abdullahelwalid 3 months ago

Selected Answer: D

Option A is not the answer because if the application is migrated to lambda then code refactoring is required which will require operational overhead, while option D the architecture remains the same but we evenly distribute the traffic by adding ALB then assigning the EC2s to a target group therefore the load will be evenly balanced. Route 53 gets updated pointing to the ALB
upvoted 1 times

ParamD 3 months, 2 weeks ago

Selected Answer: C

D. doesn't have auto scaling.
B. EKS will add operational overhead
A. Adds lots of lambda functions whose maintenance and management will add to operational overhead compared to current monolithic setup.
C. Is the best fit of the available options, it will enable autoscaling and will allow upto 8 nodes from current 5, one lambda function to update route53 will add minimal operational overhead. Though D with Autoscaling would have allowed minimal operational overhead and more flexibility to scale.
upvoted 1 times

soulation 4 months ago

Selected Answer: C

Less operational overhead. Much less development effort.
upvoted 1 times

SaqibTaqi 4 months, 3 weeks ago

Selected Answer: A

well... i have to say... none of the options here comply to least operational overhead... each and every option involves changing the application logic.. but foe the sake of it... A is the best answer.. It cannot be B as containerizing would not be suitable to use with IP addresses of the instances... ASG and ELB would not fit here as Route 53 records point to the static IP addresses of the instances.. so the best answer is A... But again... a lot of overhead invovled if someone goes on for implementation...
upvoted 1 times

sintesi_suffisso0 5 months, 1 week ago

Selected Answer: D

It can't be A since we don't know how much time the API needs to complete

upvoted 2 times

Shanmahi 5 months, 2 weeks ago

Selected Answer: D

While all 4 options work well and general inclination is to go for "serverless", the least operational effort is certainly add an ALB to distribute the incoming traffic on the EC2 instances. In a "real-world" scenario, I would ideally place Route53 -> ALB -> EC2 instances in an ASG. However, in the given option choices, D with ALB meets the requirement well from operational complexity point of view.

upvoted 3 times

jerry00218 6 months ago

Selected Answer: A

Serverless is the least operational effort

upvoted 1 times

thanhpolimi 6 months ago

Selected Answer: D

D provides a balanced solution to handle increased and varying traffic loads while minimizing the complexity and maintenance overhead.

upvoted 2 times

grumpysloth 6 months, 2 weeks ago

Selected Answer: C

Operational overhead to fix the scalabiltiy issue is minimal if we keep the EC2 instances as they are and use ASG. We know nothing about the code complexity or response time, it might be hours, so Lambda is nto an option IMHO. D is not an option because it doesn't include autoscaling, so it won't solve the issue.

upvoted 3 times

JOJO9 6 months, 3 weeks ago

Selected Answer: D

This approach leverages AWS managed services like the Application Load Balancer (ALB) and Auto Scaling groups, minimizing the operational overhead required to handle varying traffic loads. The ALB automatically distributes incoming traffic across the EC2 instances, while the instances can be placed in private subnets for better security. Additionally, the Auto Scaling group can be configured to automatically scale the EC2 instances based on metrics like CPU utilization, eliminating the need for manual scaling.

By using these managed services, you can offload tasks like load balancing, health checks, and auto-scaling to AWS, reducing the operational burden on your team. Updating the Route 53 record to point to the ALB's DNS name ensures that traffic is seamlessly routed to the backend instances without the need for manual DNS updates or additional components like Lambda functions.

upvoted 2 times

Question #34 *Topic 1*

A company has created an OU in AWS Organizations for each of its engineering teams. Each OU owns multiple AWS accounts. The organization has hundreds of AWS accounts.
A solutions architect must design a solution so that each OU can view a breakdown of usage costs across its AWS accounts.
Which solution meets these requirements?

A. Create an AWS Cost and Usage Report (CUR) for each OU by using AWS Resource Access Manager. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

C. Create an AWS Cost and Usage Report (CUR) in each AWS Organizations member account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

D. Create an AWS Cost and Usage Report (CUR) by using AWS Systems Manager. Allow each team to visualize the CUR through Systems Manager OpsCenter dashboards.

Suggested Answer: *B*

Community vote distribution

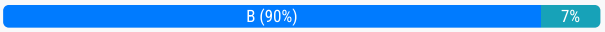

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

B is the correct answer. The solution would be to create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account. This would allow the management account to view the usage costs across all the member accounts, and the teams can visualize the CUR through an Amazon QuickSight dashboard. This allows the organization to have a centralized place to view the cost breakdown and the teams to access the cost breakdown in an easy way.

upvoted 20 times

85b5b55 Most Recent 5 months ago

Selected Answer: B

using OU Management Account, create a CUR using OU Management account and allow each AWS account to view through Amazon QuickSight.

upvoted 2 times

AWSum1 9 months, 3 weeks ago

Selected Answer: B

Option B: it must be done from the management accoint

upvoted 1 times

amministrazione 10 months ago

B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

upvoted 1 times

TonytheTiger 1 year, 1 month ago

Selected Answer: C

Option C: I hate this questions because you have 2 correct answers however only ONE real correct answer. I have to read the question like 20x until I understood it, the questions is asking for " solution so the EACH OU can VIEW a breakdown of usage across ITS account".
Its only asking for each OU breakdown for its members can see the usage cost and NOT the organization. Prior to Dec 2020 Option B would be correct however after its Option C:

Read the following AWS Update - https://aws.amazon.com/about-aws/whats-new/2020/12/cost-and-usage-report-now-available-to-member-linked-accounts/?pg=ln&sec=uc

upvoted 4 times

ParamD 3 months, 2 weeks ago

No, ask if for OU level report, not member level. Hence B.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account. Allow each team to visualize the CUR through an Amazon QuickSight dashboard.

upvoted 1 times

Rajarshi 1 year, 4 months ago

C

As target is to design a solution so that each OU can view a breakdown of usage costs across its AWS accounts

upvoted 1 times

acordovam 1 year, 4 months ago

Selected Answer: A

The question specifies that each OU should only view their own AWS accounts, not all accounts in the organization. While creating the solution in the management account might offer a centralized approach, it violates this crucial requirement.

upvoted 1 times

acordovam 1 year, 4 months ago

Sorry, I'm wrong, RAM can't create a Cost Report.

upvoted 3 times

abeb 1 year, 7 months ago

B From management account of each account

upvoted 1 times

daz2023 1 year, 9 months ago

AWS Resource Access Manager has nothing to do with creating CUR.

Answer B is correct. Use AWS Organization management account

upvoted 1 times

duriselvan 1 year, 10 months ago

https://aws.amazon.com/blogs/mt/visualize-and-gain-insights-into-your-aws-cost-and-usage-with-cloud-intelligence-dashboards-using-amazon-quicksight/

upvoted 1 times

NikkyDicky 2 years ago

Selected Answer: B

B by elimination

upvoted 1 times

gameoflove 2 years, 1 month ago

Selected Answer: B

B As AWS Organizations Management account is only correct option

upvoted 1 times

leehjworking 2 years, 2 months ago

Can anyone explain why A is wrong? Thank you.

upvoted 1 times

scuzzy2010 2 years, 2 months ago

AWS Resource Access Manager has nothing to do with creating CURs. It's for sharing resources with other accounts.

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: B

B. Create an AWS Cost and Usage Report (CUR) from the AWS Organizations management account.

upvoted 2 times

masetromain 2 years, 6 months ago

Selected Answer: B

https://www.examtopics.com/discussions/amazon/view/71951-exam-aws-certified-solutions-architect-professional-topic-1/

upvoted 3 times

Question #35 *Topic 1*

A company is storing data on premises on a Windows file server. The company produces 5 GB of new data daily.
The company migrated part of its Windows-based workload to AWS and needs the data to be available on a file system in the cloud. The company already has established an AWS Direct Connect connection between the on-premises network and AWS.
Which data migration strategy should the company use?

A. Use the file gateway option in AWS Storage Gateway to replace the existing Windows file server, and point the existing file share to the new file gateway.

B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.

C. Use AWS Data Pipeline to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).

D. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS).

Suggested Answer: *B*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.
D. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS) are also valid options. They both use DataSync to schedule a daily task to replicate the data between on-premises and cloud, the main difference is the type of file system in the cloud, Amazon FSx or Amazon Elastic File System (Amazon EFS).
upvoted 16 times

rbm2023 2 years, 1 month ago

EFS only support Linux FS. this is why we need to go for FSx . option B
upvoted 27 times

Karamen 1 year, 10 months ago

thanks for this explaination.
> EFS only support Linux FS. this is why we need to go for FSx . option B
upvoted 1 times

victorHugo Highly Voted 1 year, 10 months ago

Selected Answer: A

For an and b we need FSx. Data Sync is useful for a batch and is able to process large data volumes. in (a) the data is also accessible from on prem. The data volume is quite small (5 GB) per day therefore (a) is feasible. In my opinion, the key requirement is "data to be available on a file system in the cloud" and ",, migrating workloads" and I think this includes that it can be accessed from servers on prem. In addition (a) replaces only a Windows File server and not the overall windows landscape in AWS. There I vote for (a), AWS Data Sync.

See https://tutorialsdojo.com/aws-datasync-vs-storage-gateway/ for a comparison
upvoted 13 times

swadeey 1 year, 7 months ago

Correct point here is migration not daily sync and replication.
upvoted 3 times

vn_thanhtung 1 year, 9 months ago

needs the data to be available on a file system in the cloud
upvoted 3 times

0dc6cac Most Recent 1 week ago

Selected Answer: A

Nothing suggests that we want to do a daily sync in the question. A makes more sense because the files will be always available to use from the cloud. In 99% of the cases that's better, and if it's the remaining 1%, they would have surely mentioned it in the question.

upvoted 1 times

happpieee 3 months, 1 week ago

Selected Answer: B

DayaSync allows migration and continue access to both on-prem file server and FSx in a synchronised manner.

upvoted 1 times

ParamD 3 months, 2 weeks ago

Selected Answer: A

A seem to be a more efficient solution as it eliminates duplicate storage. Storage gateway has to be FSx, which is implicit in this option.

upvoted 1 times

fbukevin 6 months ago

Selected Answer: B

Comparing A and B, finally I choosed B due to the "scheduling". In A, despite file gateway could provide or combine a scheduling function, B explicitly says that operation.

upvoted 1 times

Biden 6 months ago

Selected Answer: A

Workloads are only partially migrated which means data need to be accessed simultaneously by on-prem VMs and the already migrated VMs in cloud
And we should assume that the data should be up to date all the time not just updated periodically
Finally, File Gateway could be Amazon FSX File GW
Hence clearly A.

upvoted 2 times

Tiger4Code 7 months ago

Selected Answer: B

D is wrong cos EFS support only Linux File System

upvoted 2 times

toyaji 8 months, 3 weeks ago

Selected Answer: A

Using Amazon FSx File Gateway, you can access data with low latency from on-premise and also in-cloud always. Why do you neet to batch datasync as like B?

upvoted 1 times

amministrazione 10 months ago

B. Use AWS DataSync to schedule a daily task to replicate data between the on-premises Windows file server and Amazon FSx.

upvoted 1 times

8693a49 11 months ago

Selected Answer: A

Because part of the workloads have already been migrated we need a solution that keeps the data consistent between on prem and the cloud. With DataSync files stored by systems on-prem would be visible in the cloud only the following day. This could cause data inconsistencies and business disruption. The best solution is to use a file gateway to maintain files synchronised at all times. 5GBs/day is easily transferable over DX

upvoted 2 times

gfhbox0083 11 months, 3 weeks ago

B, for sure.
Needs the data to be available on a file system in the cloud.

upvoted 1 times

mifune 1 year, 2 months ago

Selected Answer: B

Windows file server -> FSx (cristal clear)

upvoted 1 times

Vongolatt 1 year, 2 months ago

Selected Answer: B

A is not the data migration

upvoted 2 times

mav3r1ck 1 year, 3 months ago

Selected Answer: B

option B is the most suitable data migration strategy for the company. It leverages AWS DataSync to automate the replication of daily data increments from the on-premises Windows file server to Amazon FSx for Windows File Server. This approach provides a seamless integration for Windows-based workloads with minimal disruption and supports the company's needs for a cloud-native file system that is fully managed and integrates well with AWS services.

upvoted 1 times

8693a49 11 months ago

Batch sync is not seamles and might not be with minimal disruption depending on how it's used. Is this generated with ChatGPT?

upvoted 1 times

mav3r1ck 1 year, 3 months ago

Selected Answer: B

This option is particularly suitable for the company's requirements because it allows for scheduled daily tasks to efficiently replicate the 5 GB of new data to Amazon FSx, providing a cloud-native file system that integrates well with Windows-based workloads.

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B is the answer

upvoted 1 times

Question #36 *Topic 1*

A company's solutions architect is reviewing a web application that runs on AWS. The application references static assets in an Amazon S3 bucket in the us-east-1 Region. The company needs resiliency across multiple AWS Regions. The company already has created an S3 bucket in a second Region.
Which solution will meet these requirements with the LEAST operational overhead?

A. Configure the application to write each object to both S3 buckets. Set up an Amazon Route 53 public hosted zone with a record set by using a weighted routing policy for each S3 bucket. Configure the application to reference the objects by using the Route 53 DNS name.

B. Create an AWS Lambda function to copy objects from the S3 bucket in us-east-1 to the S3 bucket in the second Region. Invoke the Lambda function each time an object is written to the S3 bucket in us-east-1. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.

C. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.

D. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. If failover is required, update the application code to load S3 objects from the S3 bucket in the second Region.

Suggested Answer: *D*

Community vote distribution

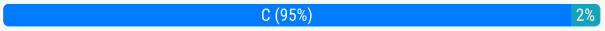

zhangyu20000 Highly Voted 2 years, 6 months ago

C is correct.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
upvoted 19 times

amministrazione Most Recent 10 months ago

C. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. Set up an Amazon CloudFront distribution with an origin group that contains the two S3 buckets as origins.
upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: C

Why not D ? D. Configure replication on the S3 bucket in us-east-1 to replicate objects to the S3 bucket in the second Region. If failover is required, update the application code to load S3 objects from the S3 bucket in the second Region:

Manual Failover: This option involves manual updates to the application code in the event of a failover, which adds operational overhead and complexity. CloudFront provides automatic failover and load balancing, making it a more streamlined solution.
upvoted 1 times

sarlos 1 year, 2 months ago

C IS THE answer
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C is correct
upvoted 1 times

VerRi 1 year, 4 months ago

Selected Answer: C

Straightforward
upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: C

Option C is the most efficient solution because it leverages S3's built-in replication feature to automatically replicate objects to a second bucket in another Region, ensuring that the data is resiliently stored across multiple Regions. By using Amazon CloudFront with an origin group containing both

S3 buckets, the application benefits from CloudFront's global content delivery network, which improves load times and provides a built-in failover mechanism. This setup minimizes operational overhead while achieving the desired resiliency and performance improvements.
Option C provides a seamless, automated solution for achieving resiliency across multiple AWS Regions with minimal operational effort, leveraging AWS services designed for replication, content delivery, and failover.
upvoted 1 times

Vaibs099 1 year, 5 months ago
C is correct because,

You can server Dynamic Websites with Static Content with CDN by having origins for both and in your webserver app refer to DNS for s3 origin from CF to deliver static content. For webserver on EC2 (Custom Origins can be used).
So in above scenario, if you would like to have resiliency. Add another S3 Origin with bucket in different region. Create Origin Group with both S3 Origins. Set priority on Origins and select 4XX and 5XX error codes for failover. You can use DNS returned for Origin Group from Cloud front in your web app and that would do automatic failover with least overheads.

D also solves the purpose, but you will need to build failover mechanism in your app. However, with above Cloudfront Origin group is taking care of that for you.
upvoted 1 times

ninomfr64 1 year, 6 months ago
Selected Answer: C
All options does the job, but:
A would require code maintenance and managing public hosted zone -> No
B would require Lambda and CloudFront operations -> No
C would require only CloudFront operations -> Yes
D requires a lot of work for failover that appears to be manual -> No
upvoted 2 times

subbupro 1 year, 6 months ago
C is mostly correct, A is not correct - B and D required the code changes. C will take care of the cloud front orgin failover.
upvoted 1 times

abeb 1 year, 7 months ago
C is good
upvoted 1 times

severlight 1 year, 7 months ago
Selected Answer: C
obvious
upvoted 1 times

totten 1 year, 9 months ago
Selected Answer: C
Here's why Option C is the most suitable choice:

Replication: Amazon S3 Cross-Region replication is designed to replicate objects from one S3 bucket to another in a different Region. This ensures data resiliency across Regions with minimal operational overhead. Once configured, replication happens automatically.

CloudFront: Setting up an Amazon CloudFront distribution with an origin group containing the two S3 buckets allows you to use a single CloudFront distribution to serve content from both Regions. CloudFront provides low-latency access to your content, and using an origin group allows for failover if one of the S3 buckets becomes unavailable.
upvoted 4 times

totten 1 year, 9 months ago
Option A suggests configuring the application to write each object to both S3 buckets, which can result in higher operational overhead and may not provide immediate failover capabilities.

Option B involves creating a Lambda function to copy objects, which adds complexity and requires code maintenance for each object written to the S3 bucket in us-east-1.

Option D relies on manual updates to the application code for failover, which is less automated and could result in higher operational overhead.

Therefore, Option C is the most efficient and operationally streamlined solution to achieve data resiliency and availability across multiple AWS Regions.

upvoted 1 times

Simon523 1 year, 9 months ago

Selected Answer: C

C, LEAST operational overhead

upvoted 1 times

TWOCATS 1 year, 10 months ago

Selected Answer: C

C should incur the least operational cost while D still requires the cx to update the code in whatever way they deem as appropriate

upvoted 1 times

Karamen 1 year, 10 months ago

Selected Answer: C

upvoted 1 times

xplusfb 1 year, 10 months ago

Selected Answer: C

Its completely asking CRR Right one is C

upvoted 1 times

Question #37 *Topic 1*

A company is hosting a three-tier web application in an on-premises environment. Due to a recent surge in traffic that resulted in downtime and a significant financial impact, company management has ordered that the application be moved to AWS. The application is written in .NET and has a dependency on a MySQL database. A solutions architect must design a scalable and highly available solution to meet the demand of 200,000 daily users.
Which steps should the solutions architect take to design an appropriate solution?

A. Use AWS Elastic Beanstalk to create a new application with a web server environment and an Amazon RDS MySQL Multi-AZ DB instance. The environment should launch a Network Load Balancer (NLB) in front of an Amazon EC2 Auto Scaling group in multiple Availability Zones. Use an Amazon Route 53 alias record to route traffic from the company's domain to the NLB.

B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

C. Use AWS Elastic Beanstalk to create an automatically scaling web server environment that spans two separate Regions with an Application Load Balancer (ALB) in each Region. Create a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a cross-Region read replica. Use Amazon Route 53 with a geoproximity routing policy to route traffic between the two Regions.

D. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon ECS cluster of Spot instances spanning three Availability Zones. The stack should launch an Amazon RDS MySQL DB instance with a Snapshot deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

Suggested Answer: *C*

Community vote distribution

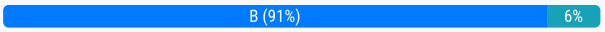

robertohyena Highly Voted 2 years, 6 months ago

Selected Answer: B

Agree with B.
Not A: we will not use NLB for web app
Not C: Beanstalk is region service. It CANNOT "automatically scaling web server environment that spans two separate Regions"
Not D: spot instances cant meet 'highly available'
upvoted 28 times

kz407 1 year, 3 months ago

I don't think ASGs are cross-region either. This answer in SO gives a serious perspective on this regard.
https://stackoverflow.com/a/12907101/3126973
upvoted 1 times

masetromain 2 years, 5 months ago

That's correct, option C is not a valid solution because AWS Elastic Beanstalk is a region-specific service, it cannot span multiple regions. Option B is a valid solution that uses CloudFormation to launch a stack with an Application Load Balancer in front of an Auto Scaling group, a Multi-AZ Aurora MySQL cluster and Route 53 to route traffic to the load balancer, it meets the requirements of scalability and high availability with a good performance and with less operational overhead.
upvoted 6 times

Perkuns 2 years ago

if I am not mistaken you can deploy the same EB to a different region. why does that eliminate C? it further increases your availability with geolocation weighted routing, as well as you having DR which even further increases availability along with low RPO and RTO
upvoted 6 times

jpa8300 1 year, 6 months ago

I agree with you, that's the best option, two EBs, one in each region to deploy, manage and monitor all the environment.
upvoted 1 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

B is correct. The solution architect should use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB

cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

This solution provides scalability and high availability for the web application by using an Application Load Balancer and an Auto Scaling group in multiple availability zones, which can automatically scale in and out based on traffic demand. The use of a Multi-AZ Amazon Aurora MySQL DB cluster provides high availability for the database layer and the Retain deletion policy ensures the data is retained even if the DB instance is deleted. Additionally, the use of Route 53 with an alias record ensures traffic is routed to the correct location.

upvoted 8 times

Deztroyer88 Most Recent 3 months, 3 weeks ago

Selected Answer: B

B is the most appropriate solution thats meeting the scalable and highly available requirement. C is good for DR and that is not what is being asked in the question.

upvoted 1 times

Archer1974 3 months, 4 weeks ago

Selected Answer: C

Highly Available is a key requirement and hence the solution has to span multiple regions. The other options do not handle for Region Failure. Elastic Beanstalk can be deployed across multiple regions. Traffic can be distributed to the appropriate region based on geo-proximity. This is the only correct answer.

upvoted 1 times

Tiger4Code 7 months ago

Selected Answer: B

Answer: B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

NOT C: Cos geoproximity is used only for private hosted zones, not public

upvoted 1 times

amministrazione 10 months ago

B. Use AWS CloudFormation to launch a stack containing an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones. The stack should launch a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy. Use an Amazon Route 53 alias record to route traffic from the company's domain to the ALB.

upvoted 1 times

gfhbox0083 11 months, 3 weeks ago

Selected Answer: B

B, for sure.
Elastic Beanstalk is region specific.
The "Retain" deletion policy in AWS Aurora ensures that when you delete a database cluster, the automated backups and snapshots of the cluster are retained. This means that even though the database cluster itself is deleted, the backups and snapshots remain, allowing you to restore the cluster from those backups at a later time.

upvoted 3 times

gfhbox0083 11 months, 3 weeks ago

B, for sure.
Elastic Beanstalk environments are typically created within a single AWS region.

upvoted 1 times

TonytheTiger 1 year, 2 months ago

Selected Answer: C

Option C: The only AWS documentation I found that support .NET application migration is for Elastic Beanstalk, it said " EB is the fastest and simplest way to deploy .NET applications on AWS" Many suggestion is selection option "B", the question is not asking about cost or least operational overhead, just scalable and highly available for the migration for a .NET application. Also, I can see why so many people are selecting option "B".

https://docs.aws.amazon.com/whitepapers/latest/develop-deploy-dotnet-apps-on-aws/aws-elastic-beanstalk.html
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/concepts.concepts.design.html

upvoted 4 times

kz407 1 year, 3 months ago

Selected Answer: B

B however is not a highly available solution IMO because it is restricted to a region. By any chance if the region goes down, the webapp goes down as well.
A is out of the picture because it involves an NLB.
D is out of the picture because it involves spot instances which is not the choice for HA requirements.
C, everything is good except the mention of "Elastic Beanstalk environment that spans across regions". This is wrong. EB environments are a region construct. You can's have them spanning cross region. You can however have EB in multiple regions.
upvoted 3 times

bjexamprep 1 year, 5 months ago

Selected Answer: B

Guessing the question designer prefers B. But it is wrong. When talking about R53 Alias record, it is wrong. Cause Alias record points to IP address while ALB endpoint is not an IP address.
A has flaw. The question says 3-tier web application. AWS question designers often mess up the definition of 3-tier application, which means there isn't a very clear definition of 3 tier: browser/application server/database is one definition, another one is WebServer/Application Server/database. Looks like A means the latter. Then, if the Elastic Beanstalk is hosting a web server, what are the ASG hosting? And why the R53 is pointing to the NLB which is pointing to the ASG?
C is wrong, cause Elastic Beanstalk cannot span regions.
D is wrong because spot instance is not HA.
Weighting the flaws of different answers, B has the least flaw.
upvoted 1 times

ninomfr64 1 year, 6 months ago

Selected Answer: B

Not C as we do not need to span multiple Region (DR, global reach, ...), also cross-Region read replica does not fail-over automatically (you need to promote it to primary). Finally from the wording it seems that this imply having a single environment that spans two separate Regions which is not supported (you need two separate environments)
Not D as we have a single RDS DB instance, no HA

Both A and B does the job, but B provides better scalability as it make use of Aurora Multi-AZ that allows secondary (reader) instance(s) to be accessed for reads, while RDS Multi-AZ instance does not allow standby instance endpoint to be accessed. This could be circumvented by using Multi-AZ DB cluster deployment that provides 2 readable standby instance
upvoted 1 times

ayadmawla 1 year, 6 months ago

Selected Answer: C

Answer is C
The best way to migrate a .NET application to AWS is via Beanstalk (see: https://docs.aws.amazon.com/whitepapers/latest/develop-deploy-dotnet-apps-on-aws/aws-elastic-beanstalk.html)

I think that the question regarding spanning a deployment across two regions has triggered some to reject based on the multi-region but if you continue you will notice the separate regional deployments based on two ALBs etc. Just my two pennie :)
upvoted 2 times

subbupro 1 year, 6 months ago

B is the correct,
upvoted 1 times

shaaam80 1 year, 6 months ago

Selected Answer: B

Answer B
upvoted 1 times

abeb 1 year, 7 months ago

B is good
upvoted 1 times

totten 1 year, 9 months ago

Selected Answer: B

Here's why Option B is the best choice:

High Availability: The use of an Application Load Balancer (ALB) and Amazon Aurora Multi-AZ deployment ensures high availability and fault tolerance for the web application and the MySQL database. The Multi-AZ setup for Aurora provides automatic failover.

Scalability: Using an EC2 Auto Scaling group across multiple Availability Zones allows the application to automatically scale to meet traffic demands. This is crucial for handling the surge in traffic from 200,000 daily users.

Deletion Policy: The Retain deletion policy for the Aurora MySQL DB cluster ensures that even if the CloudFormation stack is deleted, the database is retained, which is important for data preservation and recovery.

Route 53 Routing: Route 53 with an alias record provides efficient DNS routing, directing traffic to the ALB, which then distributes it to the EC2 instances. This ensures that users can access the application reliably.

upvoted 1 times

totten 1 year, 9 months ago

Option C introduces unnecessary complexity by spanning two separate Regions and using geoproximity routing. This is typically used for disaster recovery and global deployments, which may not be necessary here.

upvoted 1 times

Question #38 *Topic 1*

A company is using AWS Organizations to manage multiple AWS accounts. For security purposes, the company requires the creation of an Amazon Simple Notification Service (Amazon SNS) topic that enables integration with a third-party alerting system in all the Organizations member accounts.
A solutions architect used an AWS CloudFormation template to create the SNS topic and stack sets to automate the deployment of CloudFormation stacks. Trusted access has been enabled in Organizations.
What should the solutions architect do to deploy the CloudFormation StackSets in all AWS accounts?

A. Create a stack set in the Organizations member accounts. Use service-managed permissions. Set deployment options to deploy to an organization. Use CloudFormation StackSets drift detection.

B. Create stacks in the Organizations member accounts. Use self-service permissions. Set deployment options to deploy to an organization. Enable the CloudFormation StackSets automatic deployment.

C. Create a stack set in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.

D. Create stacks in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets drift detection.

Suggested Answer: *C -*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The best solution is C, because it involves creating the stack set in the management account of the organization, which is the central point of control for all the member accounts. This allows the solutions architect to manage the deployment of the stack set across all member accounts from a single location. Service-managed permissions are used, which allows the CloudFormation service to deploy the stack set to all member accounts. The deployment options are set to deploy to the organization and automatic deployment is enabled, which ensures that the stack set is automatically deployed to all member accounts as soon as it is created in the management account.

upvoted 21 times

masetromain Highly Voted 2 years, 6 months ago

Selected Answer: C

https://www.examtopics.com/discussions/amazon/view/47723-exam-aws-certified-solutions-architect-professional-topic-1/

upvoted 5 times

amministrazione Most Recent 10 months ago

C. Create a stack set in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.

upvoted 1 times

Vaibs099 1 year, 5 months ago

C. Create a stack set in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets automatic deployment.

C is more suitable as Enable CloudFormation StackSets automatic deployment will take care of any new account in the Org. Set deployment options to deploy to the organization helps deploying Stack Instances to targeted account in Org. Use service-managed permissions is hassle free as it takes care or roles for you.

D. Create stacks in the Organizations management account. Use service-managed permissions. Set deployment options to deploy to the organization. Enable CloudFormation StackSets drift detection.

D is good option too as StackSets drift detection is a good option to have but not a requirement. It only saves from future troubleshooting of drift scenarios.

upvoted 1 times

nharaz 1 year, 5 months ago

Selected Answer: C

D is wrong - Drift Detection identifies unmanaged changes (Outside CloudFormation)

upvoted 2 times

jainparag1 1 year, 7 months ago

Selected Answer: C

I'll go with C since it satisfies all the requirements with minimum operational overhead. But wondering if "Stack Sets drift detection" is just a distractor here. Can someone throw some light on this?

upvoted 2 times

ninomfr64 1 year, 6 months ago

I am not an expert, just sharing my thoughts:
"Stack Sets drift detection" is a feature of stack set, however this is not needed according to the scenario.
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacksets-drift.html.

D is a no-go for me because it deploys in each managed account without making use of stack sets, so you cannot then use stack sets drift detection.

upvoted 1 times

daz2023 1 year, 9 months ago

Selected Answer: C

C is the right answer

upvoted 1 times

NikkyDicky 2 years ago

Selected Answer: C

C no brainer

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

Create a stack set in the Organizations management account.

upvoted 2 times

spd 2 years, 4 months ago

Selected Answer: C

Stack Set in Mgmt account

upvoted 2 times

Atila50 2 years, 6 months ago

I THINK I SHOULD BE A

upvoted 1 times

Question #39 *Topic 1*

A company wants to migrate its workloads from on premises to AWS. The workloads run on Linux and Windows. The company has a large on-premises infrastructure that consists of physical machines and VMs that host numerous applications.

The company must capture details about the system configuration, system performance, running processes, and network connections of its on-premises workloads. The company also must divide the on-premises applications into groups for AWS migrations. The company needs recommendations for Amazon EC2 instance types so that the company can run its workloads on AWS in the most cost-effective manner.

Which combination of steps should a solutions architect take to meet these requirements? (Choose three.)

A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.

B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs.

C. Group servers into applications for migration by using AWS Systems Manager Application Manager.

D. Group servers into applications for migration by using AWS Migration Hub.

E. Generate recommended instance types and associated costs by using AWS Migration Hub.

F. Import data about server sizes into AWS Trusted Advisor. Follow the recommendations for cost optimization.

Suggested Answer: *BDE*

Community vote distribution

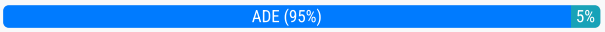

bititan Highly Voted 2 years, 5 months ago

Selected Answer: ADE

trusted advisor doesn't have option to upload data, so option F is irrelavent

upvoted 24 times

ninomfr64 Highly Voted 1 year, 6 months ago

Selected Answer: ADE

A vs B -> A because we need to use AWS Application Discovery and it provides its own agent

https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html

C vs D -> D because AWS Application Discovery is integrated with AWS Migration Hub and it can be used to group servers into applications

https://aws.amazon.com/migration-hub/faqs/#:~:text=How%20do%20I%20group%20servers%20into%20an%20application%3F

E vs. F -> E as AWS Migration Hub allows to generate recommendation for instance types

https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html

upvoted 7 times

amministrazione Most Recent 10 months ago

A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs.

D. Group servers into applications for migration by using AWS Migration Hub.

E. Generate recommended instance types and associated costs by using AWS Migration Hub.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: ADE

Why not B ? B. Assess the existing applications by installing AWS Systems Manager Agent on the physical machines and VMs:

Explanation: AWS Systems Manager Agent is used for managing and automating tasks on EC2 instances, not for capturing detailed application and performance data during an assessment phase. AWS Application Discovery Agent is more appropriate for this purpose.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: ADE

ADE is correct

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: ADE

The correct answers are:

* A. Assess the existing applications by installing AWS Application Discovery Agent on the physical machines and VMs. The AWS Application Discovery Service helps gather detailed information about on-premises data centers, including servers, network dependencies, and performance metrics.

* D. Group servers into applications for migration by using AWS Migration Hub. AWS Migration Hub provides a centralized location to track the progress of application migrations across multiple AWS and partner solutions. It allows grouping discovered servers into applications, which simplifies the organization of migration tasks.

* E. Generate recommended instance types and associated costs by using AWS Migration Hub. After servers are discovered and grouped into applications, AWS Migration Hub can analyze the collected data to recommend suitable Amazon EC2 instance types. This ensures that the migrated applications are hosted on the most cost-effective resources.

upvoted 3 times

Simon523 1 year, 9 months ago

Selected Answer: ADE

https://aws.amazon.com/tw/blogs/mt/using-aws-migration-hub-network-visualization-to-overcome-application-and-server-dependency-challenges/

upvoted 2 times

NikkyDicky 2 years ago

Selected Answer: ADE

ADE no brainer

upvoted 1 times

ZK000001qws 2 years ago

B in incorrect as System Manager doesn't do discovery however, SSM Agent makes it possible for Systems Manager to update, manage, and configure the resources in AWS as well as on-premises. ADE

upvoted 3 times

asifjanjua88 2 years, 2 months ago

ADE is correct answer.

upvoted 1 times

Jacky_exam 2 years, 2 months ago

Selected Answer: ADE

https://docs.aws.amazon.com/application-discovery/latest/userguide/discovery-agent.html

https://docs.aws.amazon.com/migrationhub/latest/ug/ec2-recommendations.html

upvoted 2 times

hgc2023 2 years, 3 months ago

B is incorrect because the servers are on prem.

upvoted 1 times

ninomfr64 1 year, 6 months ago

SSM can be installed on on-premise server. This is not the point for not picking B

upvoted 1 times

dev112233xx 2 years, 3 months ago

Selected Answer: ADE

ADE no doubts ✍

upvoted 1 times

God_Is_Love 2 years, 4 months ago

Logical answer : Falls under the domain "Accelerate Workload Migration and Modernization"

promoting MigrationHub

Step 1 - Identify the apps

Step 2 - Group them

Step 3 - Before hand, find out what instance types would need to be in when

actual migration happens

https://d1.awsstatic.com/Product-Page-Diagram_AWS-Migration-Hub-Orchestrator%402x.0c34c9483d13ebd26cf9072193384a58531624f3.png

For OnPremises migrations, first phase is Discovery which can be done with

Discovery agent , A

https://d1.awsstatic.com/products/application-discovery-service/Product-Page-Diagram_AWS-Application-Discovery-Service%201.9d81c27f3de50349a9406b8def61b8eb914e2930.png

I wont go with Trusted Advisor although it advises how cost can be advised because-
This applies for already aws available environment. Here, about to get migrated into
AWS and Architects need to discover lot of info before hand to plan alot. So I choose E between E and F. My answer - A,D,E
upvoted 2 times

aws0909 2 years, 4 months ago
Why Option C Group servers into applications for migration by using AWS Systems Manager Application Manager is incorrect?
upvoted 1 times

sambb 2 years, 4 months ago
AWS SSM Application Manager is used for existing resources deployed to AWS
upvoted 1 times

moota 2 years, 4 months ago
Selected Answer: ADE
A is better than B.

> Agent-based discovery can be performed by deploying the AWS Application Discovery Agent on each of your VMs and physical servers. The agent installer is available for Windows and Linux operating systems. It collects static configuration data, detailed time-series system-performance information, inbound and outbound network connections, and processes that are running.

https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html
upvoted 1 times

boomx 2 years, 5 months ago
BDE. Trusted Advisor is not for onprem assessments. Migration hub does EC2 ones
upvoted 1 times

Question #40 *Topic 1*

A company is hosting an image-processing service on AWS in a VPC. The VPC extends across two Availability Zones. Each Availability Zone contains one public subnet and one private subnet.

The service runs on Amazon EC2 instances in the private subnets. An Application Load Balancer in the public subnets is in front of the service. The service needs to communicate with the internet and does so through two NAT gateways. The service uses Amazon S3 for image storage. The EC2 instances retrieve approximately 1 TB of data from an S3 bucket each day.

The company has promoted the service as highly secure. A solutions architect must reduce cloud expenditures as much as possible without compromising the service's security posture or increasing the time spent on ongoing operations.

Which solution will meet these requirements?

A. Replace the NAT gateways with NAT instances. In the VPC route table, create a route from the private subnets to the NAT instances.

B. Move the EC2 instances to the public subnets. Remove the NAT gateways.

C. Set up an S3 gateway VPC endpoint in the VPAttach an endpoint policy to the endpoint to allow the required actions on the S3 bucket.

D. Attach an Amazon Elastic File System (Amazon EFS) volume to the EC2 instances. Host the images on the EFS volume.

Suggested Answer: *C*

Community vote distribution

C (100%)

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: C

C. Setting up an S3 gateway VPC endpoint in the VPC and attaching an endpoint policy to the endpoint will allow the EC2 instances to securely access the S3 bucket for image storage without the need for NAT gateways, reducing costs without compromising security or increasing ongoing operations. This option reduces the costs associated with the NAT gateways and allows for faster data retrieval from the S3 bucket as traffic does not have to go through the internet gateway.

upvoted 15 times

God_Is_Love Highly Voted 1 year, 10 months ago

The only reason for C is - Gateway endpoints are not Billed and so cost effective (https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc-endpoints-for-s3) If the question changes from single region to across region, the answer would be B (overhead of NAT gateways and traversing TBs of data across NAT is expensive) because gateway endpoints are region specific

upvoted 7 times

anita_student 1 year, 10 months ago

B wouldn't be highly secure and data transfer would also be slower

upvoted 1 times

8608f25 Most Recent 10 months, 3 weeks ago

Selected Answer: C

Option C is the most cost-effective solution that maintains the service's security posture. An S3 gateway VPC endpoint allows private connections between the VPC and S3 without requiring traffic to go through the internet or NAT gateways. This eliminates the need for NAT gateways when accessing S3, which can significantly reduce costs, especially considering the 1 TB of data retrieved daily from S3. Endpoint policies ensure that the security posture is not compromised by allowing only the required actions on the specific S3 bucket.

upvoted 1 times

grire974 11 months, 3 weeks ago

Any chance someone could fix the typo in the correct answer; "VPC. Attach..." instead of VPAttach; terribly misleading.

upvoted 2 times

daz2023 1 year, 2 months ago

Selected Answer: C

C for using an endpoint.

upvoted 2 times

NikkyDicky 1 year, 6 months ago

C of course

upvoted 1 times

gameoflove 1 year, 7 months ago

Selected Answer: C

C is the Correct option as S3 Gateway will reduce the cost for NAT gateway

upvoted 2 times

mfsec 1 year, 9 months ago

Selected Answer: C

Set up an S3 gateway VPC endpoint

upvoted 3 times

dev112233xx 1 year, 9 months ago

Selected Answer: C

C - easy one ✅

upvoted 3 times

zozza2023 1 year, 11 months ago

Selected Answer: C

C for sure

upvoted 4 times

Question #41 *Topic 1*

A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load for the rest of the week. The access pattern includes many more writes to the table than reads of the table.

A solutions architect needs to implement a solution to minimize the cost of the table.

Which solution will meet these requirements?

A. Use AWS Application Auto Scaling to increase capacity during the peak period. Purchase reserved RCUs and WCUs to match the average load.

B. Configure on-demand capacity mode for the table.

C. Configure DynamoDB Accelerator (DAX) in front of the table. Reduce the provisioned read capacity to match the new peak load on the table.

D. Configure DynamoDB Accelerator (DAX) in front of the table. Configure on-demand capacity mode for the table.

Suggested Answer: *D*

Community vote distribution

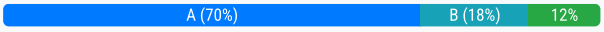

zhangyu20000 Highly Voted 2 years, 5 months ago

A is correct. On demand mode is for unknown load pattern, auto scaling is for know burst pattern

upvoted 25 times

AimarLeo 1 year, 4 months ago

But the pattern here is known.. 4 hours peak time etc.. not sure if that would be the write answer

upvoted 1 times

dqwsmwwvtgxwkvgcvc 1 year, 10 months ago

How AWS Application Auto Scaling scale the read/write performance of DynamoDB?

upvoted 1 times

tannh 1 year, 9 months ago

You can scale DynamoDB tables and global secondary indexes using target tracking scaling policies and scheduled scaling.
https://docs.aws.amazon.com/autoscaling/application/userguide/services-that-can-integrate-dynamodb.html

upvoted 1 times

ccort Highly Voted 2 years, 5 months ago

Selected Answer: A

A

on-demand prices can be 7 times higher, given the options it is better to have reserved WCU and RCU and auto scale in the given schedule

upvoted 16 times

zhen234 Most Recent 5 months, 2 weeks ago

Selected Answer: A

Reserved capacity applies to the baseline level of provisioned throughput. During peak workloads, Auto Scaling dynamically adjusts the provisioned capacity of a DynamoDB table (RCUs and WCUs) based on the actual workload. you are charged on-demand rates for the excess capacity.

upvoted 1 times

wem 6 months, 3 weeks ago

Selected Answer: B

B. Configure on-demand capacity mode for the table.

Explanation:
On-Demand Capacity Mode:

DynamoDB's on-demand capacity mode automatically adjusts to accommodate variable workloads.

It eliminates the need to provision RCUs and WCUs, allowing the table to scale up during the 4-hour peak period and scale down during off-peak times, which is cost-effective when usage is highly variable.
Cost Optimization:

With on-demand capacity, you pay only for the read and write requests that are made. This is ideal for workloads with sporadic or unpredictable traffic patterns, such as this scenario with a weekly 4-hour peak.
Minimized Operational Overhead:

On-demand mode requires no manual adjustments or additional services (like Application Auto Scaling), simplifying management and reducing costs related to provisioning errors or overprovisioning.
Access Pattern with More Writes:

On-demand capacity mode is well-suited for write-heavy workloads as it scales automatically to handle higher write throughput during peak times.
upvoted 3 times

5e8c031 3 days, 3 hours ago
Except that this scenario with a known average and a weekly 4-hour peak is neither unpredictable nor sporadic.
upvoted 1 times

Sin_Dan 8 months, 2 weeks ago
Selected Answer: B
B is the right answer. Reserved RCU/WCU locks you into fixed cost. Even though on demand is more expensive, the additional cost is paid only for 4 hrs a week.
upvoted 2 times

DhirajBansal 7 months ago
but Yes, here in option A it is saying for purchasing average load RCUs and WCUs which will cost less and also auto scaling can be used for scheduled scaling WCU and RCUs which will be cost efficient.
upvoted 1 times

amministrazione 10 months ago
A. Use AWS Application Auto Scaling to increase capacity during the peak period. Purchase reserved RCUs and WCUs to match the average load.
upvoted 1 times

subbupro 10 months, 1 week ago
I think B is correct. because reserved is not required, ondemand would be better because it requireds only 4 hours per week. so B would be better. Autoscaling of the application can not impact dynamo db tables.
upvoted 1 times

vn_hunglv 11 months, 2 weeks ago
Selected Answer: A
Tôi chọn A
upvoted 1 times

zolthar_z 11 months, 2 weeks ago
Selected Answer: A
Auto-scaling is for known traffic pattern, On-demand is for unknown traffic patter and also could be more expensive
upvoted 2 times

Malcnorth59 1 year, 1 month ago
Selected Answer: A
AWS documentation suggests A is correct:
https://docs.aws.amazon.com/autoscaling/application/userguide/what-is-application-auto-scaling.html
upvoted 3 times

mnsait 7 months, 1 week ago
Nice. Thanks for the link. It explains clearly.

See this: "Scheduled scaling – Scale a resource one time only or on a recurring schedule."
upvoted 1 times

Kubernetes 1 year, 2 months ago
A is correct. The focus is minimizing the cost of tables.
upvoted 2 times

mav3r1ck 1 year, 3 months ago

Selected Answer: B

Considering the application's need to handle a peak load that is double the average and the fact that the workload is write-heavy, option B (Configure on-demand capacity mode for the table) is the most suitable solution. It directly addresses the variability in workload without requiring upfront capacity planning or additional management overhead, thus likely providing the best cost optimization for this scenario. On-demand capacity mode eliminates the need to scale resources manually or through Auto Scaling and ensures that you only pay for the write and read throughput you consume.

upvoted 2 times

mav3r1ck 1 year, 3 months ago

A. AWS Application Auto Scaling with Reserved Capacity

Pros: Auto Scaling allows you to automatically adjust the provisioned throughput to meet demand, and purchasing reserved RCUs and WCUs can reduce costs for the capacity you know you'll consistently use.

Cons: This option might not be as cost-effective for workloads with significant variability and a high write-to-read ratio, especially if the peak load is much higher than the average load. Reserved capacity benefits consistent usage patterns, but the peak load being double the average may not be fully optimized here.

upvoted 1 times

mav3r1ck 1 year, 3 months ago

B. On-demand Capacity Mode

Pros: On-demand capacity mode is ideal for unpredictable workloads because it automatically scales to accommodate the load without provisioning. You pay for what you use without managing capacity planning. This mode is particularly suitable for the described scenario where the load spikes significantly and unpredictably.

Cons: While potentially more expensive per unit than provisioned capacity with auto-scaling, it eliminates the risk of over-provisioning or under-provisioning.

upvoted 1 times

kz407 1 year, 3 months ago

Selected Answer: A

A is badly worded however, because it says "application" autoscaling. We are not talking about that here. Either it should be reworded as "DynamoDB autoscaling" for the answer to be correct.

On-demand capacity mode is for unknown read/write patterns. Since the load change patterns are known, anything that involves on-demand capacity modes can be eliminated (hence not B).

DAX is a caching service deployed in front of DynamoDB. It is geared towards "performance at scale". Problem in the use case, is to optimize table costs. Using DAX will incur additional costs. Hence anything that involves DAX (C and D) can also be eliminated.

upvoted 2 times

Malcnorth59 1 year, 1 month ago

I initially thought the same but the AWS definition of Application autoscaling listed here includes DynamoDB:

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

upvoted 1 times

anubha.agrahari 1 year, 3 months ago

Selected Answer: A

https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization-at-any-scale/#:~:text=You%20can%20approximate%20a%20blend,save%20money%20as%20reserved%20capacity

upvoted 2 times

8608f25 1 year, 4 months ago

Selected Answer: B

Option B is the most cost-effective solution for workloads with significant fluctuations and unpredictable access patterns. The on-demand capacity mode automatically adjusts the table's throughput capacity as needed in response to actual traffic, eliminating the need to manually configure or manage capacity. This mode is ideal for applications with irregular traffic patterns, such as a significant peak once a week, because you only pay for the read and write requests your application performs, without having to provision throughput in advance. Option B directly addresses the requirement to minimize costs associated with fluctuating loads, especially when the load significantly exceeds the average only during a brief period, by leveraging DynamoDB's on-demand capacity mode to automatically scale and pay only for what is used.

upvoted 1 times

igor12ghsj577 1 year, 4 months ago

Selected Answer: A

I think there is mistake in answer A, and it should be DynamoDb auto scaling instead of application autos calling. Or application and dynamoDB auto scaling.

upvoted 1 times

igor12ghsj577 1 year, 4 months ago

Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf, in response to actual traffic patterns. This enables a table or a global secondary index to increase its provisioned read and write capacity to handle sudden increases in traffic, without throttling. When the workload decreases, Application Auto Scaling decreases the throughput so that you don't pay for unused provisioned capacity.

upvoted 2 times

jpa8300 1 year, 6 months ago

Selected Answer: D

I choose option D, because DAX is not only an accelerator for the Reads, it also cache releasing a lot of load from the DB.

upvoted 1 times

longlehoang 4 months, 2 weeks ago

The access pattern includes many more writes to the table than read. so I think D incorrect

upvoted 1 times

Question #42 *Topic 1*

A solutions architect needs to advise a company on how to migrate its on-premises data processing application to the AWS Cloud. Currently, users upload input files through a web portal. The web server then stores the uploaded files on NAS and messages the processing server over a message queue. Each media file can take up to 1 hour to process. The company has determined that the number of media files awaiting processing is significantly higher during business hours, with the number of files rapidly declining after business hours.

What is the MOST cost-effective migration recommendation?

A. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in an Amazon S3 bucket.

B. Create a queue using Amazon MQ. Configure the existing web server to publish to the new queue. When there are messages in the queue, create a new Amazon EC2 instance to pull requests from the queue and process the files. Store the processed files in Amazon EFS. Shut down the EC2 instance after the task is complete.

C. Create a queue using Amazon MQ. Configure the existing web server to publish to the new queue. When there are messages in the queue, invoke an AWS Lambda function to pull requests from the queue and process the files. Store the processed files in Amazon EFS.

D. Create a queue using Amazon SQS. Configure the existing web server to publish to the new queue. Use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the files. Scale the EC2 instances based on the SQS queue length. Store the processed files in an Amazon S3 bucket.

Suggested Answer: *D*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: D

The correct answer would be option D.

This option suggests creating a queue using Amazon SQS, configuring the existing web server to publish to the new queue, and using EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the files. The EC2 instances can be scaled based on the SQS queue length, which ensures that the resources are available during peak usage times and reduces costs during non-peak times.

Option A is not correct because it suggests using AWS Lambda which has a maximum execution time of 15 minutes.
Option B is not correct because it suggests creating a new EC2 instance for each message in the queue, which is not cost-effective.
Option C is not correct because it suggests using Amazon EFS, which is not a suitable option for long-term storage of large files.

upvoted 24 times

ninomfr64 Highly Voted 1 year, 6 months ago

Selected Answer: D

Not A - Lambda max execution time is 15 minutes, image processing can take up to 1 hour
Not B - Amazon MQ is not needed (more expensive then SQS) and EFS is more expensive then S3
Not C - Amazon MQ is not needed (more expensive then SQS) and Lambda max execution time is 15 minutes, image processing can take up to 1 hour

D does the job with the lower cost thanks to SQS, S3 and EC2 Auto Scaling Group

upvoted 7 times

mohan_cv Most Recent 8 months ago

Exam AWS Certified Solutions Architect - Professional SAP-C02 topic 1 question 42 discussion

upvoted 1 times

Malcnorth59 1 year, 1 month ago

Selected Answer: D

Lambda will not work, so A is not possible.
D is going to be the most cost-effective as the resources will scale based on queue length.

upvoted 1 times

mav3r1ck 1 year, 3 months ago

Selected Answer: D

Given the need to process files that can take up to 1 hour each and the variability in workload, option D (Amazon SQS, EC2 Auto Scaling, and S3) appears to be the most cost-effective and practical solution. It leverages SQS for queue management, enabling efficient handling of the processing queue's variability. EC2 Auto Scaling allows for flexible and cost-effective scaling of processing capacity, ramping up during high-demand periods and scaling down when demand wanes, thus optimizing costs. Finally, Amazon S3 offers a highly durable and cost-effective solution for storing the processed media files. This option provides the necessary flexibility for long processing tasks while efficiently managing the variable demand and optimizing storage costs.

upvoted 1 times

Simon523 1 year, 9 months ago

Selected Answer: D

Simple Queuing Service

SQS is based on pull model. Here are some of the important features:

Reliable, scalable, fully-managed message queuing service
High availability
Unlimited scaling
Auto scale to process billions of messages per day
Low cost (Pay for use)

upvoted 1 times

aviathor 1 year, 10 months ago

Selected Answer: D

This is quite simple. Any answer (A and C) consisting of using Lambda for processing the files is out because of the 15 minutes limit on Lambda processes.

B is out because using EFS is expensive and it does not specify how to launch and terminate the EC2 instances. Amazon MQ is not required either.

This leaves D which uses SQS, Auto Scaling Groups and publishes the resulting files to S3.

upvoted 2 times

chico2023 1 year, 10 months ago

Selected Answer: D

Answer: D

You can eliminate A and C right in the beginning: Lambda functions can run up to 15 minutes.
B won't help much as you need to create new EC2 instances (manually, apparently) and EFS is more expensive than S3.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

d for sure

upvoted 1 times

ailves 2 years ago

Selected Answer: D

Because of "Each media file can take up to 1 hour to process" and we know Lambda has a limit in 15 minutes, The correct answer is D

upvoted 1 times

EricZhang 2 years, 1 month ago

D - https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-using-sqs-queue.html

upvoted 1 times

huanaws088 2 years, 2 months ago

Selected Answer: B

I sure is B , becauce

1. SQS , SNS are " cloud - native " services : proprietary protocols from AWS
2. Traditional applications running from on - premises may use open protocols such as : MQTT , AMQP ,.., so When migrating to the cloud , instead of re-engineering the application to use SQS and SNS will very expensive, we can use Amazon MQ.
3. Amazon MQ doesn't " scale " as much as SQS / SNS Amazon MQ runs on servers but Amazon MQ has both queue feature (~ SQS) and topic features (~ SNS)

https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-difference-from-amazon-mq-sns.html
upvoted 1 times

hexie 1 year, 11 months ago
In terms of cost (which is a point on the question), Amazon SQS is generally more cost-effective compared to Amazon MQ for this specific use case. SQS pricing is based on the number of requests and message data transfer, whereas Amazon MQ pricing includes additional costs associated with broker instances and data transfer.
upvoted 1 times

takecoffe 2 years, 2 months ago
Selected Answer: D
SQS and autoscaling no doubt answer is D
upvoted 2 times

mfsec 2 years, 3 months ago
Selected Answer: D
SQS and Auto Scaling
upvoted 2 times

dev112233xx 2 years, 3 months ago
Selected Answer: D
D - makes sense.. Lambda can't run more than 15m.
And Amazon MQ is only recommended when migrating existing message brokers that rely on compatibility with APIs such as JMS or protocols such as AMQP, MQTT, OpenWire, and STOMP.. in the question there is no mention for these services ..
upvoted 4 times

God_Is_Love 2 years, 4 months ago
A and C are out because lambda does not support more than 15 min. B says, to create an EC2 for each new message which is certainly not cost effective and bad design as well. So answer is D
upvoted 2 times

c73bf38 2 years, 4 months ago
Selected Answer: D
The most cost-effective migration recommendation to handle peak loads during business hours is to use Amazon SQS to create a queue, configure the existing web server to publish to the new queue, and use Amazon EC2 instances in an EC2 Auto Scaling group to pull requests from the queue and process the files. The EC2 instances should be scaled based on the SQS queue length. Storing the processed files in an Amazon S3 bucket will help in reducing the storage cost. This approach is scalable and can handle peak loads during business hours, while still being cost-effective during non-business hours. Option A is also a possible solution, but using EC2 instances in an EC2 Auto Scaling group is a more scalable and cost-effective solution. Options B and C involve using Amazon EFS, which can be more expensive than Amazon S3.
upvoted 2 times

Question #43 *Topic 1*

A company is using Amazon OpenSearch Service to analyze data. The company loads data into an OpenSearch Service cluster with 10 data nodes from an Amazon S3 bucket that uses S3 Standard storage. The data resides in the cluster for 1 month for read-only analysis. After 1 month, the company deletes the index that contains the data from the cluster. For compliance purposes, the company must retain a copy of all input data.

The company is concerned about ongoing costs and asks a solutions architect to recommend a new solution.

Which solution will meet these requirements MOST cost-effectively?

A. Replace all the data nodes with UltraWarm nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

B. Reduce the number of data nodes in the cluster to 2 Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Transition the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy.

C. Reduce the number of data nodes in the cluster to 2. Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Add cold storage nodes to the cluster Transition the indexes from UltraWarm to cold storage. Delete the input data from the S3 bucket after 1 month by using an S3 Lifecycle policy.

D. Reduce the number of data nodes in the cluster to 2. Add instance-backed data nodes to handle the expected capacity. Transition the input data from S3 Standard to S3 Glacier Deep Archive when the company loads the data into the cluster.

Suggested Answer: *B*

Community vote distribution

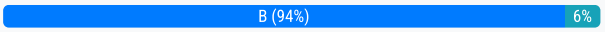

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

B is the most cost-effective solution as it reduces the number of data nodes in the cluster to 2 and adds UltraWarm nodes to handle the expected capacity. By configuring the indexes to transition to UltraWarm when OpenSearch Service ingests the data, the company can take advantage of the lower storage costs of UltraWarm. Additionally, by transitioning the input data to S3 Glacier Deep Archive after 1 month using an S3 Lifecycle policy, the company can further reduce costs by using the lower storage costs of S3 Glacier Deep Archive for long-term data retention.

upvoted 21 times

masetromain 2 years, 5 months ago

Option C can meet the requirements of reducing the number of data nodes in the cluster and using UltraWarm and cold storage nodes to handle the expected capacity and moving the data to lower cost storage after 1 month. However, it may not be the most cost-effective solution as it involves additional complexity in configuring the indexes to transition between different storage tiers, and may also require additional management and maintenance of the cold storage nodes. Option B, where the data is transitioned from S3 Standard to S3 Glacier Deep Archive using an S3 Lifecycle policy is simpler and more cost-effective as it eliminates the need for additional storage tiers and management.

upvoted 3 times

God_Is_Love 2 years, 4 months ago

B says to delete but question asks for saving on compliance purposes.

upvoted 5 times

God_Is_Love 2 years, 4 months ago

* I meant C says..

upvoted 5 times

amministrazione Most Recent 10 months ago

B. Reduce the number of data nodes in the cluster to 2 Add UltraWarm nodes to handle the expected capacity. Configure the indexes to transition to UltraWarm when OpenSearch Service ingests the data. Transition the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy.

upvoted 1 times

Malcnorth59 1 year, 1 month ago

Why can't I switch all nodes to ultrawarm. I can't find it anywhere in the documentation and it's not listed in the pre-requisites.

Also why can the number of nodes be reduced from 10 to 2? is that because Ultrawarm use S3?
upvoted 1 times

sarlos 1 year, 2 months ago
why not D?
upvoted 1 times

ninomfr64 1 year, 6 months ago
I need help here:

To use UltraWarm storage, domains must have dedicated master nodes as per doc https://docs.aws.amazon.com/opensearch-service/latest/developerguide/ultrawarm.html

The scenario mentions "an OpenSearch Service cluster with 10 data nodes". Assuming you only have these nodes in the cluster, in all answers you need to add dedicated master node(s). Assuming we also have dedicated master node why not replacing all data nodes with UltraWarm nodes?
upvoted 1 times

ninomfr64 1 year, 6 months ago
I think I got it, UltraWarm is for read-only data. Thus you still need to have at least a data node
upvoted 1 times

venvig 1 year, 10 months ago
Selected Answer: B
Option A says to replace all Data Nodes with ultra warm nodes. But this is NOT possible. There has to be atleast one data node
upvoted 3 times

NikkyDicky 1 year, 12 months ago
Selected Answer: B
B I think :/
upvoted 2 times

Damijo 2 years, 3 months ago
Selected Answer: A
If you look at the IAM documentation here, you can see that the ec2:AuthorizeSecurityGroupIngress action doesn't have any conditions that would allow you to specify the ip addresses in the inbound/outbound rules.https://docs.aws.amazon.com/service-authorization/latest/reference/list_amazonec2.html
upvoted 2 times

Jesuisleon 2 years ago
I think you are referring All AWS Certified Solutions Architect - Professional SAP-C02 Questions, question 44. yes, I changed from D to A after reading this link.
upvoted 1 times

eddylynx 1 year, 12 months ago
You can specify the IP address with the CIDR parameter

https://ec2.amazonaws.com/?Action=AuthorizeSecurityGroupIngress
&GroupId=sg-112233
&IpPermissions.1.IpProtocol=tcp
&IpPermissions.1.FromPort=3389
&IpPermissions.1.ToPort=3389
&IpPermissions.1.IpRanges.1.CidrIp=192.0.2.0/24
&IpPermissions.1.IpRanges.1.Description=Access from New York office

https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_AuthorizeSecurityGroupIngress.html
upvoted 1 times

dev112233xx 2 years, 3 months ago
Selected Answer: B
B - makes more sense
upvoted 4 times

Ajani 2 years, 3 months ago

UltraWarm provides a cost-effective way to store large amounts of read-only data on Amazon OpenSearch Service. Standard data nodes use "hot" storage, which takes the form of instance stores or Amazon EBS volumes attached to each node. Hot storage provides the fastest possible performance for indexing and searching new data.

upvoted 3 times

moota 2 years, 4 months ago

I asked ChatGPT. Can I use all UltraWarm nodes in AWS OpenSearch instead of data nodes? :)

No, UltraWarm nodes in AWS OpenSearch are designed for storage and retrieval of infrequently accessed data, while data nodes are optimized for faster indexing and searching of data. While UltraWarm nodes can be used as a complement to data nodes, they are not a replacement for them.

upvoted 2 times

hobokabobo 2 years, 4 months ago

This eliminates option A

upvoted 2 times

Musk 2 years, 5 months ago

Selected Answer: B

Option B is the most cost-effective solution that meets the requirements. Reducing the number of data nodes in the cluster and adding UltraWarm nodes will help to reduce the ongoing costs of running the OpenSearch Service cluster. Configuring the indexes to transition to UltraWarm when OpenSearch Service ingests the data will further reduce costs. Additionally, transitioning the input data to S3 Glacier Deep Archive after 1 month by using an S3 Lifecycle policy will lower the storage costs of retaining the input data for compliance purposes.

upvoted 4 times

Question #44 *Topic 1*

A company has 10 accounts that are part of an organization in AWS Organizations. AWS Config is configured in each account. All accounts belong to either the Prod OU or the NonProd OU.

The company has set up an Amazon EventBridge rule in each AWS account to notify an Amazon Simple Notification Service (Amazon SNS) topic when an Amazon EC2 security group inbound rule is created with 0.0.0.0/0 as the source. The company's security team is subscribed to the SNS topic.

For all accounts in the NonProd OU, the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source.

Which solution will meet this requirement with the LEAST operational overhead?

A. Modify the EventBridge rule to invoke an AWS Lambda function to remove the security group inbound rule and to publish to the SNS topic. Deploy the updated rule to the NonProd OU.

B. Add the vpc-sg-open-only-to-authorized-ports AWS Config managed rule to the NonProd OU.

C. Configure an SCP to allow the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is not 0.0.0.0/0. Apply the SCP to the NonProd OU.

D. Configure an SCP to deny the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. Apply the SCP to the NonProd OU.

Suggested Answer: *C*

Community vote distribution

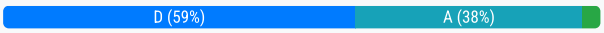

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: D

The solution that meets this requirement with the LEAST operational overhead is D. Configuring an SCP to deny the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0, and applying the SCP to the NonProd OU. This solution would prevent the security group inbound rule from being created in the first place and will not require any additional steps or actions to be taken in order to remove the rule. This is less operationally intensive than modifying the EventBridge rule to invoke an AWS Lambda function, adding a Config rule or allowing the ec2:AuthorizeSecurityGroupIngress action with a specific IP.

upvoted 52 times

masetromain 2 years, 5 months ago

Option C does not meet the requirement that the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source. It only allows the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is not 0.0.0.0/0. It does not prevent the creation of a security group inbound rule that includes 0.0.0.0/0 as the source, it only allows for the ingress action on non-0.0.0.0/0 IPs.

Option D is the best solution as it denies the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. This will prevent the creation of any security group inbound rule that includes 0.0.0.0/0 as the source.

upvoted 6 times

MikelH93 2 years, 1 month ago

the answer can't be C or D because aws:SourceIp condition key don't exist with SCP.

So answer is A

upvoted 4 times

mifune 1 year, 2 months ago

You mean something like this? It's from the AWS portal...

{

"Version": "2012-10-17",

"Statement": {

"Effect": "Deny",

"Action": "*",

"Resource": "*",
"Condition": {
"NotIpAddress": {
"aws:SourceIp": [
"192.0.2.0/24",
"203.0.113.0/24"
]
}
}
}
}
upvoted 4 times

b3llman 1 year, 10 months ago
have you actually tested it? if you haven't, please do it and then comment.
upvoted 4 times

aokaddaoc 1 year, 7 months ago
I think the reason why C is wrong is not because C does not meet the requirement but simply because it is too strong: All users can do is to set ingress rule in SG and all other actions are all blocked. Both C and D results the same which users can no longer able to open port to 0.0.0.0/0, but D is more precise without blocking other actions.
upvoted 2 times

Maria2023 Highly Voted 2 years ago
Selected Answer: D
I literally just created the SCP and it works. I saw some comments that "ec2:AuthorizeSecurityGroupIngress action doesn't have any conditions" - that is not correct. This is my scp :
{
"Sid": "Statement1",
"Effect": "Deny",
"Action": [
"ec2:AuthorizeSecurityGroupIngress"
],
"Resource": [
"*"
],
"Condition": {
"IpAddress": {
"aws:SourceIp": [
"0.0.0.0/0"
]
}
}
}
upvoted 35 times

tgv 5 months ago
Tested myself, but this blocks any attempt to create an ingress rule - not only ones that have 0.0.0.0/0 as a source.

aws:SourceIp checks for the IP address of the requester
https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition-keys.html#condition-keys-sourceip

With these options, I think the only option that still stands is [A[.
I don't like it because it adds management overhead and it's not preventive - it's reactive. But it seems like the only one which actually performs the task it was asked to perform.
upvoted 1 times

12db8b7 4 days, 10 hours ago
A can't be because u can create the resource and then is deleted, the question ask to "needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source", in A scenario u can do it, just deletes the rule afterwards.
upvoted 1 times

b3llman 1 year, 10 months ago

Tested and confirmed!

upvoted 6 times

dqwsmwwvtgxwkvgcvc 1 year, 10 months ago

I guess proving D works doesn't show C is incorrect. I feel that both C and D could be correct because as CuteRunRun mentioned, the SCP deny is default.

Just have one more question, what is the ec2:AuthorizeSecurityGroupIngress if the SourceIp is not 0.0.0.0/0?

upvoted 1 times

vn_thanhtung 1 year, 10 months ago

For all accounts in the NonProd OU, the security team needs to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source.

you think C can "remove the ability to create" carry ? SCP allow all by default?

upvoted 1 times

vn_thanhtung 1 year, 10 months ago

Sorry typo.

you think C can "remove the ability to create" crazy ? SCP allow all by default

upvoted 1 times

longns 1 year, 9 months ago

This will deny all action create a inbound rule not only Inbound rule which have source ip "0.0.0.0/0"

upvoted 4 times

Malcnorth59 1 year, 1 month ago

I think that is incorrect. the SCP action is ec2:AuthorizeSecurityGroupIngress and specifically applies to ingress

upvoted 1 times

proawsk Most Recent 2 months ago

Selected Answer: B

B is correct, you cannot deny SG rule creation with SCP

upvoted 1 times

tgv 5 months ago

Selected Answer: A

aws:SourceIp checks for the ip address of the requester - not the CIDR destination in the rule

upvoted 1 times

TorTo 5 months ago

Selected Answer: D

The only correct answer is D.

The questions states "to remove the ability to create a security group inbound rule that includes 0.0.0.0/0 as the source"

A does not remove the ability, it only corrects the action.

D is correct because it actually restricts the ability.

upvoted 2 times

altonh 5 months, 3 weeks ago

Selected Answer: A

Not D. See here: https://docs.aws.amazon.com/service-authorization/latest/reference/list_amazonec2.html. Condition key aws:SourceIp is missing for ec2:AuthorizeSecurityGroupIngress

upvoted 1 times

grumpysloth 6 months, 2 weeks ago

Selected Answer: A

You cannot use SCP to control SG rules

upvoted 1 times

chipimbiri 7 months ago

Selected Answer: D

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_evaluation.html

An Allow statement in an SCP can't have a Condition element at all.

upvoted 1 times

sashenka 9 months ago

Selected Answer: A

Given that the aws:SourceIp condition key refers to the IP address of the principal making the request, and not the IP address specified in the security group rule, D is not appropriate for this scenario.

upvoted 3 times

attila9778 6 months, 4 weeks ago

But because of this B is the correct option.

upvoted 1 times

amministrazione 10 months ago

D. Configure an SCP to deny the ec2:AuthorizeSecurityGroupIngress action when the value of the aws:SourceIp condition key is 0.0.0.0/0. Apply the SCP to the NonProd OU.

upvoted 1 times

[Removed] 10 months ago

Selected Answer: D

Service Control Policy (SCP):

Restrictive Policy Enforcement: An SCP (Service Control Policy) is used in AWS Organizations to enforce account-level restrictions across accounts that belong to a particular Organizational Unit (OU). By setting an SCP to deny the ec2:AuthorizeSecurityGroupIngress action when the aws:SourceIp condition is 0.0.0.0/0, you effectively prevent all users within the NonProd OU from creating any security group rule that opens inbound traffic to the entire internet.
Least Operational Overhead: SCPs are centrally managed and enforced automatically, requiring no further intervention once applied. This reduces the operational overhead to nearly zero, as it does not require ongoing monitoring, function deployments, or manual rule updates.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: D

Why Option D is Better than Option C:
Explicit Deny vs. Implicit Allow:

Option C allows the action unless the aws:SourceIp is 0.0.0.0/0. This creates an implicit allow policy, which means that if any condition is not met, the action is allowed.
Option D uses an explicit deny, which is more secure and straightforward. An explicit deny ensures that if the condition is met (aws:SourceIp is 0.0.0.0/0), the action is blocked regardless of other permissions.

upvoted 3 times

asquared16 11 months, 3 weeks ago

Selected Answer: A

It's A. Definitely A. Don't get confused.

upvoted 1 times

dzidis 1 year ago

Voting for A

upvoted 1 times

teo2157 1 year, 1 month ago

Selected Answer: A

It's A, D is incorrect as it shouldn´t be source IP but destination address

upvoted 1 times

Malcnorth59 1 year, 1 month ago

Selected Answer: D

Option D

upvoted 1 times

sse69 1 year, 1 month ago

Selected Answer: A

SourceIP is for requester IP address, not the CIDR referenced in the SG rule.

upvoted 3 times

Question #45 *Topic 1*

A company hosts a Git repository in an on-premises data center. The company uses webhooks to invoke functionality that runs in the AWS Cloud. The company hosts the webhook logic on a set of Amazon EC2 instances in an Auto Scaling group that the company set as a target for an Application Load Balancer (ALB). The Git server calls the ALB for the configured webhooks. The company wants to move the solution to a serverless architecture.

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

A. For each webhook, create and configure an AWS Lambda function URL. Update the Git servers to call the individual Lambda function URLs.

B. Create an Amazon API Gateway HTTP API. Implement each webhook logic in a separate AWS Lambda function. Update the Git servers to call the API Gateway endpoint.

C. Deploy the webhook logic to AWS App Runner. Create an ALB, and set App Runner as the target. Update the Git servers to call the ALB endpoint.

D. Containerize the webhook logic. Create an Amazon Elastic Container Service (Amazon ECS) cluster, and run the webhook logic in AWS Fargate. Create an Amazon API Gateway REST API, and set Fargate as the target. Update the Git servers to call the API Gateway endpoint.

Suggested Answer: *C*

Community vote distribution

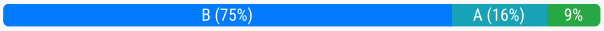

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

B. Create an Amazon API Gateway HTTP API. Implement each webhook logic in a separate AWS Lambda function. Update the Git servers to call the API Gateway endpoint. This solution will provide low operational overhead as it utilizes the serverless capabilities of AWS Lambda and API Gateway, which automatically scales and manages the underlying infrastructure and resources. It also allows for the webhook logic to be easily managed and updated through the API Gateway interface.

The answer should be B because it is the best solution in terms of operational overhead.

upvoted 24 times

masetromain 2 years, 5 months ago

Option A would require updating the Git servers to call individual Lambda function URLs for each webhook, which would be more complex and time-consuming than calling a single API Gateway endpoint.

Option C would require deploying the webhook logic to AWS App Runner, which would also be more complex and time-consuming than using an API Gateway.

Option D would also require containerizing the webhook logic and creating an ECS cluster and Fargate, which would also add complexity and operational overhead compared to using an API Gateway.

upvoted 8 times

hobokabobo 2 years, 4 months ago

I do agree with B.

However on Git server side it does make no difference if one calls aws or do a rest call via gateway.
Eg. if you use Python it makes no difference if you use boto(call Lambda) or request(rest api) module.
If one implemets via shell it makes no difference if one uses aws-cli(invoke Lambda directly) or curl(do a rest call).
Similar for other implementations.

upvoted 2 times

hobokabobo 2 years, 4 months ago

As addition why B is still better: it hides the implementation details and decouples by introducing a interface.
With that a team for Aws may change what ever it needs to change to implement the interface. On the other hand on git side can use whatever deems necessary without caring about implementation details.

upvoted 2 times

ninomfr64 Highly Voted 1 year, 6 months ago

Selected Answer: A

I need help here: what's wrong with Lambda Function URL?

With A I just need to handle my Lambda functions, updates go trough updating my aliases pointing to a new version. Here I am just missing all the capabilities provided by API Gateway that seems not to be requested (transformations, throttling, quotas, cache, api keys, auth, OpenAPI, ...). With B I still need to implement each webhook logic in a separate AWS Lambda function and update git server + I need to operates API Gateway.

Any other option requires 2 or more services thus generating more operations, also:
Not C as app runner is not a target for ALB (private IP, ECS, EC2 instance, Lambda)
Not D as you cannot set Fargate as API Gateway target (while you can use ECS as target)

Can you help me understand why B requires less operations overhead?

upvoted 6 times

Malcnorth59 1 year, 1 month ago

Option A requires that you update the webhooks for each lambda function. This will create a considerable operational overhead not just for the initial change but going forward as well.

API Gateway (B) decouple the functions from the Webhooks.

upvoted 2 times

glf Most Recent 5 months, 2 weeks ago

Selected Answer: A

B is plain wrong. With an HTTP API Gateway you're publishing your API on the Internet, which could be forbidden by security policy. Even if it is not, you at least need to setup authentication, with OAuth2 or custom Lambda authorizer. And in this case, you lose the advantage in terms of operational overhead.

upvoted 1 times

GabrielShiao 9 months ago

Selected Answer: C

Deploy the web hook logic to the Apprunner which takes a minor effort to build and deploy the container image automatically without underlying infrastructure management.

upvoted 1 times

amministrazione 10 months ago

B. Create an Amazon API Gateway HTTP API. Implement each webhook logic in a separate AWS Lambda function. Update the Git servers to call the API Gateway endpoint.

upvoted 1 times

subbupro 10 months, 1 week ago

C is the best one . Operational over head - exisiting web logic needs to be change into the lambda. But in C - just we can use the same logic just deployment activities. Please go with C

upvoted 1 times

Malcnorth59 1 year, 1 month ago

Selected Answer: B

A: large operational overhead
B: Choice
C: App runner doesn't use ALB
D: Unnecessary complexity with containers

upvoted 3 times

Fu7ed 1 year, 1 month ago

https://aws.amazon.com/ko/solutions/implementations/git-to-s3-using-webhooks/

upvoted 2 times

Fu7ed 1 year, 1 month ago

choose B

upvoted 1 times

kz407 1 year, 3 months ago

Selected Answer: B

Given the current answers, I think B is the only possible option with least overhead.

C would have been a better candidate over B, if it mentioned to include the App Runner in a Target Group TG and assign TG as the target for the API Gateway. As it stands now, C is not correct because App Runner app can't be directly assigned as a target for API Gateway.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A, because is the solution with less operational overhead. Also option B also will create new lambda functions per webhook, and you have to define the specific path in the apigateway and integrate it with your specific lambda...

upvoted 5 times

bjexamprep 1 year, 3 months ago

Selected Answer: B

Lambda function is the easiest way to implement the webhook logic. App Runner and ECS all requires more ops overhead.
So the answer is between A and B. Someone argue that using A introduces ops overhead of mapping every Lambda function to the webhooks, but actually with B, users don't need to map Lambda function in git webhooks, but move the Lambda function mapping ops to API gateway. The mapping need to be done, that's an ops overhead that cannot be ignored.
I'm guessing the question designer prefers to use API GW, because the description "Update the Git servers to call the individual Lambda function URLs." doesn't look good. While, in reality, the repo developers create the Lambda function, and they know the URL, it's very easy to launch the Lambda function from the web hook. No additional API GW is required.

upvoted 1 times

master9 1 year, 5 months ago

Selected Answer: C

You can set App Runner as a target for ALB.

AWS App Runner can use your code. You can use AWS App Runner to create and manage services based on two fundamentally different service sources: source code and source image. App Runner starts, runs, scales, and balances your service regardless of the source type. You can use the CI/CD capability of App Runner to track changes to your source image or code. When App Runner discovers a change, it automatically builds (for source code) and deploys the new version to your App Runner service

upvoted 1 times

djeong95 1 year, 4 months ago

Looks like App Runner is built more for deploying web applications rather than hosting webhook logics.

upvoted 1 times

uas99 1 year, 6 months ago

A. is the right answer as no need to introduce gateway here

upvoted 2 times

subbupro 1 year, 6 months ago

Least operations is the key. App runner is a aws managed one and can deploy it easily, A and B we need to create lamda for each web hook it is very complex . So C would be correct

upvoted 1 times

jpa8300 1 year, 6 months ago

ninomfr64 says that App runner cannot be a target for ALB, so that's the reason you cannot select C.

upvoted 2 times

severlight 1 year, 7 months ago

Selected Answer: B

Don't see the exact reasons to not choose A for now, but B will work for sure.

upvoted 1 times

severlight 1 year, 7 months ago

UPD: Don't see the exact reasons why A won't work for now, but B will work for sure.

upvoted 1 times

whenthan 1 year, 8 months ago

Selected Answer: B

reducing operational overhead!

upvoted 1 times

Andy97229 1 year, 8 months ago

Selected Answer: C

B vs C. Looking at App Runner C makes more sense.

upvoted 1 times

Question #46 *Topic 1*

A company is planning to migrate 1,000 on-premises servers to AWS. The servers run on several VMware clusters in the company's data center. As part of the migration plan, the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes. The company then wants to query and analyze the data.

Which solution will meet these requirements?

A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises hosts. Configure Data Exploration in AWS Migration Hub. Use AWS Glue to perform an ETL job against the data. Query the data by using Amazon S3 Select.

B. Export only the VM performance information from the on-premises hosts. Directly import the required data into AWS Migration Hub. Update any missing information in Migration Hub. Query the data by using Amazon QuickSight.

C. Create a script to automatically gather the server information from the on-premises hosts. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hub. Query the data directly in the Migration Hub console.

D. Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3.

Suggested Answer: *C*

Community vote distribution

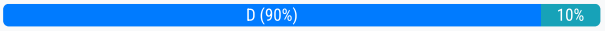

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: D

The correct answer is D: Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3.

Here is why the other choices are not correct:

A. Deploy and configure the AWS Agentless Discovery Connector virtual appliance on the on-premises hosts. Configure Data Exploration in AWS Migration Hub. Use AWS Glue to perform an ETL job against the data. Query the data by using Amazon S3 Select. - AWS Agentless Discovery Connector will help in discovering and inventory servers but it does not provide the same level of detailed metrics as the AWS Application Discovery Agent, it also does not cover process information.

upvoted 47 times

masetromain 2 years, 5 months ago

B. Export only the VM performance information from the on-premises hosts. Directly import the required data into AWS Migration Hub. Update any missing information in Migration Hub. Query the data by using Amazon QuickSight. - It does not cover process information and it's not the best way to collect the required data, it's not efficient and it might miss some important information.

C. Create a script to automatically gather the server information from the on-premises hosts. Use the AWS CLI to run the put-resource-attributes command to store the detailed server data in AWS Migration Hub. Query the data directly in the Migration Hub console. - this solution might not be very reliable and it does not cover process information, also it does not provide a way to query and analyze the data.

upvoted 6 times

masetromain 2 years, 5 months ago

D. Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3. - This is the correct answer as it covers all the requirements mentioned in the question, it will allow collecting the detailed metrics, including process information and it provides a way to query and analyze the data using Amazon Athena.

upvoted 5 times

icassp Highly Voted 2 years, 5 months ago

Selected Answer: D

Choosing between A and D. For A, how can S3 select query?

upvoted 6 times

oatif 2 years, 4 months ago

I think A is a better solution because the Agentless discovery connector is custom-made for the VMware environment. It will save us time and collect all the necessary data we need. Installing a Discovery agent in every server would be very time-consuming. S3 select allows simple select operations against your raw data. I don't think we need athena for

upvoted 4 times

djeong95 1 year, 4 months ago

As written by jainparag1, S3 Select is definitely the wrong solution here. As you said, it only allows for very simple select operations. Athena is a better way to go once you have configured the Migration hub settings correctly.

upvoted 1 times

jainparag1 1 year, 7 months ago

A is horrible. You can write only simple SQLs using S3 select. But here you need a sophisticated solution to query these special metrics. D is satisfying all the requirements.

upvoted 3 times

Jorkaef Most Recent 7 months, 3 weeks ago

A is correct

upvoted 1 times

liuliangzhou 9 months, 2 weeks ago

Selected Answer: D

If precise information about each running Process is required, it is necessary to consider using Agent-based Discovery.

upvoted 2 times

amministrazione 10 months ago

D. Deploy the AWS Application Discovery Agent to each on-premises server. Configure Data Exploration in AWS Migration Hub. Use Amazon Athena to run predefined queries against the data in Amazon S3.

upvoted 1 times

Jason666888 11 months ago

Selected Answer: D

D for sure

upvoted 1 times

vip2 1 year, 1 month ago

Selected Answer: D

see https://docs.aws.amazon.com/application-discovery/latest/userguide/what-is-appdiscovery.html

for VMs hosted on VMware, you can use both the Agentless Collector and Discovery Agent to perform discovery simultaneously.

Agentless Collector captures system performance information and resource utilization for each VM running in the vCenter, regardless of what operating system is in use. However, it cannot “look inside” each of the VMs, and as such, cannot figure out what processes are running on each VM nor what network connections exist. Therefore, if you need this level of detail and want to take a closer look at some of your existing VMs in order to assist in planning your migration, you can install the Discovery Agent on an as-needed basis.

upvoted 3 times

gofavad926 1 year, 3 months ago

Selected Answer: D

D is correct

upvoted 1 times

whichonce 1 year, 4 months ago

Selected Answer: A

Definetely A

https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-collector-data-collected-vmware.html

Vmware supports agentless connector with AWS, and data can be imported ove Migration Hub

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: D

Option D is the most efficient and streamlined solution for the requirements. Deploying the AWS Application Discovery Agent on each on-premises server allows for detailed collection of server metrics, including CPU usage, RAM usage, operating system details, and running processes. By

configuring Data Exploration in AWS Migration Hub, the collected data can be analyzed and queried effectively. Using Amazon Athena for querying enables powerful SQL-based exploration of the data stored in Amazon S3, offering a flexible and scalable way to analyze the migration readiness and planning data.

It is not option C because, Option C involves creating a custom script to gather server information and using the AWS CLI to store data in AWS Migration Hub. While this approach could potentially work, it requires significant manual effort to develop, deploy, and maintain the scripts across 1,000 servers, which is not ideal for minimizing operational overhead.

upvoted 1 times

ninomfr64 1 year, 6 months ago

Selected Answer: D

Not A - as AWS Agentless Discovery Connector does not provide processes visibility
Not B - as Migration Hub Import functionality does not support process datahttps://docs.aws.amazon.com/cli/latest/reference/mgh/put-resource-attributes.html, also I do not see how to query with QuickSight as there is not direct integration with Migration Hub to my knowledge
Not C - as it seems that put-resource-attributes command does not support process data
https://docs.aws.amazon.com/cli/latest/reference/mgh/put-resource-attributes.html

D is correct as Discovery Agent collects the required data including processes, Data Exploration in Migration Hub allows to use Amazon Athena and comes with pre-defined queries as well. https://docs.aws.amazon.com/application-discovery/latest/userguide/explore-data.html

upvoted 1 times

edder 1 year, 7 months ago

Selected Answer: D

https://docs.aws.amazon.com/application-discovery/latest/userguide/explore-data.html

upvoted 1 times

punkbuster 1 year, 10 months ago

Selected Answer: D

The agent-based collector can collect data related to running processes which is not available to the Agentless Collector.

Check out for yourself in the FAQs:
https://aws.amazon.com/application-discovery/faqs/

upvoted 1 times

xplusfb 1 year, 10 months ago

Selected Answer: A

As far as i learned for VM based envs we can go with agentless. And we can use a OVA image via collect the metrics and so on. im going with A .
https://docs.aws.amazon.com/application-discovery/latest/userguide/agentless-data-collected.html

upvoted 2 times

chico2023 1 year, 10 months ago

Selected Answer: D

Answer: D

The requirement: "the company wants to gather server metrics such as CPU details, RAM usage, operating system information, and running processes."

From https://aws.amazon.com/application-discovery/faqs/:

=== AWS Application Discovery Service Discovery Agent
Q: What data does the AWS Application Discovery Service Discovery Agent capture?
The Discovery Agent captures system configuration, system performance, running processes, and details of the network connections between systems.

upvoted 1 times

chico2023 1 year, 10 months ago

=== Agentless Collector
Q: What data does the Agentless Collector capture?
The Agentless Collector is delivered as an Open Virtual Appliance (OVA) package that can be deployed to a VMware host. The type of data collected will depend on the capabilities that you configure. If the credentials are provided to connect to vCenter, the Agentless Collector will collect VM inventory, configuration, and performance history data such as CPU, memory, and disk usage. If credentials are provided to connect to databases such as Oracle, SQL Server, MySQL, or PostgreSQL, the Agentless Collector will collect version, edition, and schema data. Server and

database information is uploaded to the Application Discovery Service data store. Database information can be sent to AWS DMS Fleet Advisor for analysis.

upvoted 1 times

CuteRunRun 1 year, 11 months ago

Selected Answer: D

I prefer D

upvoted 1 times

ggrodskiy 1 year, 11 months ago

Correct A.

D uses agent-based discovery, which requires installing an agent on each on-premises server. This can be cumbersome and intrusive for a large number of servers. It also does not explain how to use AWS Glue to perform an ETL job against the data.

upvoted 1 times

Question #47 *Topic 1*

A company is building a serverless application that runs on an AWS Lambda function that is attached to a VPC. The company needs to integrate the application with a new service from an external provider. The external provider supports only requests that come from public IPv4 addresses that are in an allow list.

The company must provide a single public IP address to the external provider before the application can start using the new service.

Which solution will give the application the ability to access the new service?

A. Deploy a NAT gateway. Associate an Elastic IP address with the NAT gateway. Configure the VPC to use the NAT gateway.

B. Deploy an egress-only internet gateway. Associate an Elastic IP address with the egress-only internet gateway. Configure the elastic network interface on the Lambda function to use the egress-only internet gateway.

C. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the Lambda function to use the internet gateway.

D. Deploy an internet gateway. Associate an Elastic IP address with the internet gateway. Configure the default route in the public VPC route table to use the internet gateway.

Suggested Answer: *C*

Community vote distribution

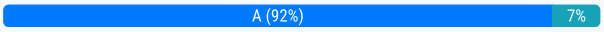

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

A. Deploy a NAT gateway. Associate an Elastic IP address with the NAT gateway. Configure the VPC to use the NAT gateway.

This solution will give the Lambda function access to the internet by routing its outbound traffic through the NAT gateway, which has a public Elastic IP address. This will allow the external provider to whitelist the single public IP address associated with the NAT gateway, and enable the application to access the new service.

upvoted 31 times

Jacky_exam 2 years, 2 months ago

Options A are not appropriate solutions because they involve deploying a NAT gateway or an egress-only internet gateway, which are used for different purposes, such as allowing resources in a private subnet to access the internet while using a static public IP address. These options will not provide the Lambda function with a single public IP address to be used for external requests.

upvoted 5 times

ninomfr64 1 year, 6 months ago

The question includes "The external provider supports only requests that come from public IPv4 addresses that are in an allow list" this imply the Lambda needs to call the external provider

upvoted 1 times

JMAN1 1 year, 6 months ago

Big Thank to you. masetromain.

upvoted 2 times

vvahe Highly Voted 2 years, 3 months ago

A

https://docs.aws.amazon.com/lambda/latest/operatorguide/networking-vpc.html

"By default, Lambda functions have access to the public internet. This is not the case after they have been configured with access to one of your VPCs. If you continue to need access to resources on the internet, set up a NAT instance or Amazon NAT Gateway. Alternatively, you can also use VPC endpoints to enable private communications between your VPC and supported AWS services."

upvoted 9 times

toyaji Most Recent 8 months, 3 weeks ago

Selected Answer: A

There are many misleading explanations here.
You cannot attath ElasticIP to Internet Gateway which use instance public ip for NAT. - https://docs.aws.amazon.com/vpc/latest/userguide/vpc-igw-internet-access.html#ip-addresses-and-nat
But NAT can be used with Elastic IP for fixed outbound ip. That's difference. - https://docs.aws.amazon.com/ko_kr/vpc/latest/userguide/nat-gateway-scenarios.html#private-nat-allowed-range
upvoted 1 times

amministrazione 10 months ago
A. Deploy a NAT gateway. Associate an Elastic IP address with the NAT gateway. Configure the VPC to use the NAT gateway.
upvoted 1 times

subbupro 10 months, 1 week ago
A is correct, NAT not only provides the internet outbound , but also provides single public IP address, So Selected Answer: A
upvoted 1 times

Jason666888 11 months ago
THE ANSWER HAS TO BE A!!!!

For B:
Wrong. Egress only internet gateway is for IPV6, not for IPV4
For C&D:
Internet gateway is for both inbound and outbount traffic. In our case we only need outbound traffic, so it has to be NAT Gateway.
upvoted 4 times

Helpnosense 1 year ago
Selected Answer: D
NAT gateway doesn't allow inbound traffic flow into service behind NAT gateway. ALB or internet gateway can. However internet gateway can't be attached to lambda service directly. I vote D as correct answer.
upvoted 2 times

kz407 1 year, 3 months ago
Selected Answer: A
Option A will be the only solution that matches the given requirements.

The problem with any solution that involves IGw is that IGw DOES NOT perform NAT. In fact, it does not alter the source IP field at all, meaning that we don't really have a mechanism of having a static public IP address set to the outbound traffic, while ensuring security. So, the only practical solution is to go with the NAT option.
upvoted 3 times

gofavad926 1 year, 3 months ago
Selected Answer: A
A, deploy nat gateway and associate an elastic ip
upvoted 1 times

Dgix 1 year, 3 months ago
Can an admin please take a look at _all_ the "correct answers" in this exam? They really cannot be trusted and reduce the usefulness of ExamTopics altogether. As things are, you should always just disregard the correct answer as it so often is insane.

The correct answer is of course A.
upvoted 3 times

Vsos_in29 1 year, 4 months ago
A is correct option, Other approach to enable internet access
https://www.linkedin.com/pulse/aws-lambda-accessing-private-vpc-resources-internet-without-vokhmin-pyxbe/
upvoted 1 times

8608f25 1 year, 4 months ago
Selected Answer: A
The solution that enables the Lambda function in a VPC to access an external service that requires requests to come from a specific public IPv4 address, and to provide a single public IP address for allow listing, is:
* Option A is correct because a NAT (Network Address Translation) gateway allows instances or AWS Lambda functions in a private subnet of a VPC to initiate outbound traffic to the internet (or external services) while preventing unsolicited inbound traffic from the internet. By associating an Elastic IP address with the NAT gateway, all outbound traffic from the Lambda function routed through the NAT gateway will appear to come from this single public IP address, which can be provided to the external provider for allow listing.

upvoted 2 times

8608f25 1 year, 4 months ago

It is not option C because, Option C describes deploying an internet gateway and associating an Elastic IP address with it. However, Lambda functions cannot be directly associated with Elastic IP addresses, and internet gateways are used to route traffic between a VPC and the internet, not to provide a static public IP address for outbound traffic.

upvoted 3 times

ninomfr64 1 year, 6 months ago

Selected Answer: A

Not B. egress-only internet gateway is IPv6 only, the question is about IPv6
Not C. you cannot associated Elastic IP to IGW also Lambda deployed in VPC cannot egress to internet via IGW, you need a NAT Gateway / NAT Instance
Not D. same as C.

A is the right solution (even if it is not well explained in my opinion)

upvoted 1 times

cgsoft 1 year, 6 months ago

Selected Answer: A

As per https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html, "To access private resources, connect your function to private subnets. If your function needs internet access, use network address translation (NAT). Connecting a function to a public subnet doesn't give it internet access or a public IP address."

upvoted 1 times

enk 1 year, 7 months ago

Selected Answer: A

Just to clarify...If the Lambda function is already attached to a VPC, it's implied that it's in a private subnet since Lambda functions can't be directly placed in public subnets. So C and D are out.

upvoted 2 times

Pupu86 1 year, 7 months ago

Selected Answer: A

Option B is definitely out as egress-only internet gateway is applicable solely for IPv6 traffic.

upvoted 2 times

whenthan 1 year, 8 months ago

Selected Answer: A

internet gateway - cant assign elastic IP to internet gateway

upvoted 1 times

Question #48 *Topic 1*

A solutions architect has developed a web application that uses an Amazon API Gateway Regional endpoint and an AWS Lambda function. The consumers of the web application are all close to the AWS Region where the application will be deployed. The Lambda function only queries an Amazon Aurora MySQL database. The solutions architect has configured the database to have three read replicas.

During testing, the application does not meet performance requirements. Under high load, the application opens a large number of database connections. The solutions architect must improve the application's performance.

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

A. Use the cluster endpoint of the Aurora database.

B. Use RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database.

C. Use the Lambda Provisioned Concurrency feature.

D. Move the code for opening the database connection in the Lambda function outside of the event handler.

E. Change the API Gateway endpoint to an edge-optimized endpoint.

Suggested Answer: *BD*

Community vote distribution

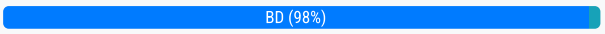

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: BD

The correct answer is B and D.

B. Using RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database can help improve the performance of the application by reducing the number of connections opened to the database. RDS Proxy manages the connection pool and routes incoming connections to the available read replicas, which can help with connection management and reduce the number of connections that need to be opened and closed.

D. Moving the code for opening the database connection in the Lambda function outside of the event handler can help to improve the performance of the application by allowing the database connection to be reused across multiple requests. This avoids the need to open and close a new connection for each request, which can be time-consuming and resource-intensive.

upvoted 48 times

masetromain 2 years, 5 months ago

A. Using the cluster endpoint of the Aurora database instead of the reader endpoint would not help improve performance in this case, because the solution architect is already using read replicas to offload read traffic from the primary instance.

C. Using the Lambda Provisioned Concurrency feature would not help improve performance in this case, as the problem is related to the number of connections to the database, not the number of instances running the Lambda function.

E. Changing the API Gateway endpoint to an edge-optimized endpoint would not help improve performance in this case, as the problem is related to the number of connections to the database, not the location of the API Gateway endpoint.

upvoted 14 times

mnsait 7 months, 1 week ago

This phrase helped me understand why A is not correct "the solution architect is already using read replicas to offload read traffic from the primary instance". Thank you @masetromain for the explanation.

upvoted 1 times

amministrazione Most Recent 10 months ago

B. Use RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database.

D. Move the code for opening the database connection in the Lambda function outside of the event handler.

upvoted 1 times

Malcnorth59 1 year, 1 month ago

Selected Answer: BD

The issue is with the number of database connections, thee are the only two changes that would impact the number of concurrent DB connections.
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: BD

B and D
upvoted 1 times

totten 1 year, 9 months ago

Selected Answer: BD

B. Use RDS Proxy to set up a connection pool to the reader endpoint of the Aurora database.

RDS Proxy helps manage and efficiently pool database connections, reducing the number of database connections required by the application. It helps improve performance and reduces the load on the database.
D. Move the code for opening the database connection in the Lambda function outside of the event handler.

By reusing database connections, you can reduce the overhead of opening and closing connections for each Lambda invocation. You can use the Lambda execution context to keep the database connection open and reuse it across multiple requests within the same execution context.
upvoted 3 times

NikkyDicky 1 year, 12 months ago

Selected Answer: BD

BD for sure
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: BD

RDS proxy + Lambda function
upvoted 4 times

dev112233xx 2 years, 3 months ago

Selected Answer: BD

RDX proxy & connecting outside the handler method is up to 5 times faster than connecting inside.
upvoted 3 times

kiran15789 2 years, 3 months ago

Selected Answer: BD

he Lambda function only queries an Amazon Aurora MySQL database- so i would reject option C
upvoted 2 times

God_Is_Love 2 years, 4 months ago

This may be too logical answer :-) - Setting up RDS proxy will help connection pooling, So B is one answer. Now C vs D
This question focuses on serverless solutions and best practices of lambda. and question hints that lambda only contains simple code.so lambda concurrency improvements may not be be the cause for performance issues detected while testing, and guess what - app is still in testing phase. so code might have a flaw can be reviewed and changed as per lambda best practices - https://docs.aws.amazon.com/lambda/latest/dg/best-practices.html. I choose B and D
upvoted 3 times

moota 2 years, 4 months ago

Selected Answer: BD

According to ChatGPT,
By reusing the same database connection across multiple invocations of the function, you can reduce the number of database connections that are opened and closed, which can help conserve resources and reduce the risk of running into database connection limits.
upvoted 2 times

Amac1979 2 years, 4 months ago

BD
https://awstut.com/en/2022/04/30/connect-to-rds-outside-of-lambda-handler-method-to-improve-performance-en/
upvoted 4 times

masssa 2 years, 5 months ago

B/C
lambda provisioned concurrency and RDS proxy are mentioned in same page.
https://quintagroup.com/blog/aws-lambda-provisioned-concurrency

upvoted 1 times

Untamables 2 years, 5 months ago

Selected Answer: BC

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.howitworks.html

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

upvoted 1 times

jhonivy 2 years, 5 months ago

B/C

Provisioned Concurrency needed: https://www.reddit.com/r/aws/comments/gcwtqt/lambda_provisioned_concurrency_with_aurora/

With connection Pool, no to worry D

upvoted 1 times

Question #49 *Topic 1*

A company is planning to host a web application on AWS and wants to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.

Which solution will meet this requirement?

A. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Export the SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.

B. Associate the EC2 instances with a target group. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure it to use the SSL certificate. Set CloudFront to use the target group as the origin server.

C. Place the EC2 instances behind an Application Load Balancer (ALB) Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Provision a third-party SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.

D. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instance. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

Suggested Answer: *C*

Community vote distribution

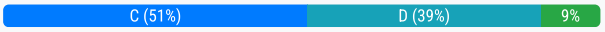

pitakk Highly Voted 2 years, 5 months ago

Selected Answer: C

Amazon-issued public certificates can't be installed on an EC2 instance. To enable end-to-end encryption, you must use a third-party SSL certificate.
https://aws.amazon.com/premiumsupport/knowledge-center/acm-ssl-certificate-ec2-elb/ so it's C or D. I choose C as it's ALB

upvoted 49 times

Jassybanga 1 year, 4 months ago

in C , the encryption will terminate at ALB so its not an end-2-end encryption , for e2e end encryption need NLB

upvoted 3 times

hobokabobo 2 years, 4 months ago

correct, but then you would use that ordered certificate for the alb as well. The other reason to order certificates is because some clients cannot verify ACM certificates which is not acceptable for a productive public service.

Between ALB and EC2 a self signed certificate is sufficient as alb does no verification of the EC2's certificate at all.

upvoted 2 times

bjexamprep 1 year, 2 months ago

that means you are decrypting the data on ALB and encrypt it again to send it to EC2. Does that sound E2E?

upvoted 5 times

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: D

Vote D.
If you need to pass encrypted traffic to targets without the load balancer decrypting it, you can create a Network Load Balancer or Classic Load Balancer with a TCP listener on port 443.
https://docs.aws.amazon.com/elasticloadbalancing/latest/application/create-https-listener.html

upvoted 41 times

hobokabobo 2 years, 4 months ago

coorect. but they want to upload the the certificate to the NLB for unknown reasons.

upvoted 7 times

Arnaud92 2 years, 3 months ago

You can use NLB with ACM cert on it. NLB can do TLS termination (https://aws.amazon.com/blogs/aws/new-tls-termination-for-network-load-balancers/) and re-encrypt to target

upvoted 2 times

lkyixoayffasdrlaqd 2 years, 4 months ago

how can this be true? Option D says to install on NLB.

You say bypass the NLB. If you bypass the NLB why are you installing the cert?

upvoted 12 times

eesa Most Recent 2 months, 1 week ago

Selected Answer: C

Why option C is correct:

End-to-end encryption implies traffic is encrypted both from the client to the load balancer and from the load balancer to the EC2 instances.

Application Load Balancers (ALBs) support HTTPS termination at the ALB level using AWS Certificate Manager (ACM) certificates.

To encrypt traffic between the ALB and EC2 instances, you must install a separate SSL/TLS certificate directly onto each EC2 instance.

AWS Certificate Manager (ACM) certificates cannot be exported or installed directly on EC2 instances; thus, a third-party SSL certificate (such as from Let's Encrypt or a commercial provider) must be used on the EC2 instances themselves.

upvoted 1 times

Trap_D0_r 4 months ago

Selected Answer: C

Please read the question carefully: "end-to-end encryption IN TRANIST"--There is no requirement for TLS termination at the NLB, and uploading the certificate to the NLB would effectively negate this anyway (I think it's thrown in there specifically to show this is the wrong answer). While it's worded poorly, the only good answer is C, which will decrypt and reencrypt traffic at the ALB only, but all traffic traversing the network will be encrypted while IN TRANSIT.

upvoted 1 times

uffd 4 months, 1 week ago

Selected Answer: D

A is not correct because as pitakk mentioned, Amazon-issued public certificates from AWS Certificate Manager (ACM) cannot be directly installed on an EC2 instance. It requires 3rd party certificates.

B doesn't make any sense.

C is not correct because ALB decrypts the traffic before sending it to the target EC2 instances.

D is correct because NLB has TCP pass through. With this, NLB doesn't have to decrypt the traffic before forwarding it to the target instances.

Courtesy to Perplexity & DeepSeek.

upvoted 1 times

attila9778 6 months, 4 weeks ago

Selected Answer: C

https://docs.aws.amazon.com/acm/latest/userguide/acm-services.html

AWS Certificate Manager (ACM) certificates cannot be directly installed on Amazon EC2 instances, except for those connected to a Nitro Enclave.

Therefore my choice is also C.

upvoted 1 times

Heman31in 6 months, 3 weeks ago

from your link : You cannot associate ACM certificates with an EC2 instance that is not connected to a Nitro Enclave. this is for Nitro case . Also : ACM is integrated with Elastic Load Balancing to deploy ACM certificates on the load balancer. For more information, so Answer is A .

upvoted 1 times

sergza 6 months, 4 weeks ago

Selected Answer: D

According to: https://docs.aws.amazon.com/elasticloadbalancing/latest/application/create-https-listener.html If you need to pass encrypted traffic to targets without the load balancer decrypting it, you can create a Network Load Balancer or Classic Load Balancer with a TCP listener on port 443. With a TCP listener, the load balancer passes encrypted traffic through to the targets without decrypting it.

upvoted 3 times

henrikhmkhitaryan59 7 months, 1 week ago

Selected Answer: D

end-to-end encryption

upvoted 3 times

AWSum1 9 months, 3 weeks ago

Selected Answer: D

I'm leaning closer to D because, NLB supports e2e. I feel that if the question asked about offloading then the ALB options may have been better. But here it's asking for e2e and can only be done with an NLB

upvoted 3 times

amministrazione 10 months ago

D. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instance. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.

upvoted 1 times

toma 11 months, 1 week ago

it is D, C is more complex.

upvoted 1 times

higashikumi 1 year, 1 month ago

Selected Answer: C

To achieve end-to-end encryption for a web application using AWS, place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM) and associate it with the ALB to handle HTTPS traffic from clients to the ALB. Additionally, install a third-party SSL certificate on each EC2 instance to ensure that traffic between the ALB and the instances is also encrypted. Configure the ALB to listen on port 443 and forward traffic to port 443 on the instances. This setup ensures that all data in transit is encrypted from the client through the ALB to the backend EC2 instances, meeting security requirements for end-to-end encryption while leveraging ACM for simplified certificate management .

upvoted 1 times

Malcnorth59 1 year, 1 month ago

Selected Answer: D

The key here is end-to-end, so that rules out ALB. Instead Use NLB with TLS termination which will pass the traffic on encrypted.

https://docs.aws.amazon.com/elasticloadbalancing/latest/network/create-tls-listener.html#:~:text=The%20load%20balancer%20passes%20the,combination%20of%20protocols%20and%20ciphers.

upvoted 2 times

titi_r 1 year, 1 month ago

Selected Answer: D

“To enable END-TO-END encryption, you must procure an SSL certificate from a third-party vendor.
You can then install the certificate on the EC2 instance and also associate the SAME certificate with the (network) Load Balancer by importing it into Amazon Certificate Manager.”
https://www.youtube.com/watch?v=6Nz0RFfBqVE&t=44s

TLS listeners for your Network Load Balancer
"… if you need to pass encrypted traffic to the targets without the (network) load balancer decrypting it, create a TCP listener on port 443 instead of creating a TLS listener."
https://docs.aws.amazon.com/elasticloadbalancing/latest/network/create-tls-listener.html

P.S. The answer is misleading because it says to install the certificate on the NLB; read it as “import it to ACM and associate it with the NLB.

upvoted 3 times

vip2 1 year, 1 month ago

Selected Answer: C

C is correct because
ALB+Self-signed Certification
NLB+Public Certification

upvoted 1 times

EmmanuelPR 1 year, 3 months ago

Selected Answer: A. Public Certificates: You can request Amazon-issued public certificates from ACM. ACM manages the renewal and deployment of public certificates that are used with ACM-integrated services, including Amazon CloudFront, Elastic Load Balancing, and Amazon API Gateway.
https://aws.amazon.com/es/certificate-manager/faqs/

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C: use ACM in the ALB and third-party SSL certificate in the EC2 instances

upvoted 2 times

Question #50 *Topic 1*

A company wants to migrate its data analytics environment from on premises to AWS. The environment consists of two simple Node.js applications. One of the applications collects sensor data and loads it into a MySQL database. The other application aggregates the data into reports. When the aggregation jobs run, some of the load jobs fail to run correctly.

The company must resolve the data loading issue. The company also needs the migration to occur without interruptions or changes for the company's customers.

What should a solutions architect do to meet these requirements?

A. Set up an Amazon Aurora MySQL database as a replication target for the on-premises database. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind a Network Load Balancer (NLB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the NLB.

B. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Move the aggregation jobs to run against the Aurora MySQL database. Set up collection endpoints behind an Application Load Balancer (ALB) as Amazon EC2 instances in an Auto Scaling group. When the databases are synced, point the collector DNS record to the ALDisable the AWS DMS sync task after the cutover from on premises to AWS.

C. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.

D. Set up an Amazon Aurora MySQL database. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as an Amazon Kinesis data stream. Use Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database. When the databases are synced, disable the replication job and restart the Aurora Replica as the primary instance. Point the collector DNS record to the Kinesis data stream.

Suggested Answer: *C*

Community vote distribution

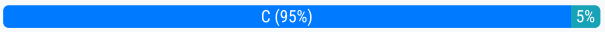

OCHT Highly Voted 2 years, 2 months ago

Selected Answer: C

Option A, B and D have some similarities with Option C but also have some key differences:

Option A uses a Network Load Balancer (NLB) instead of an Application Load Balancer (ALB) and does not use AWS Database Migration Service (AWS DMS) for continuous data replication. Instead, it sets up the Aurora MySQL database as a replication target for the on-premises database.
Option B does use AWS DMS for continuous data replication and sets up collection endpoints behind an ALB as Amazon EC2 instances in an Auto Scaling group. However, it does not create an Aurora Replica for the Aurora MySQL database or use Amazon RDS Proxy to write to the Aurora MySQL database.
Option D does not use AWS DMS for continuous data replication or set up collection endpoints behind an ALB. Instead, it sets up collection endpoints as an Amazon Kinesis data stream and uses Amazon Kinesis Data Firehose to replicate the data to the Aurora MySQL database.

upvoted 19 times

amministrazione Most Recent 10 months ago

C. Set up an Amazon Aurora MySQL database. Use AWS Database Migration Service (AWS DMS) to perform continuous data replication from the on-premises database to Aurora. Create an Aurora Replica for the Aurora MySQL database, and move the aggregation jobs to run against the Aurora Replica. Set up collection endpoints as AWS Lambda functions behind an Application Load Balancer (ALB), and use Amazon RDS Proxy to write to the Aurora MySQL database. When the databases are synced, point the collector DNS record to the ALB. Disable the AWS DMS sync task after the cutover from on premises to AWS.

upvoted 1 times

ninomfr64 1 year, 6 months ago

Selected Answer: C

Not A. not clear how the on-premises database is replicated on the Aurora MySQL, also you cannot place Lambda behind NLB as BLB only supports private IPs, instances and ALB https://docs.aws.amazon.com/elasticloadbalancing/latest/network/load-balancer-target-groups.html
Not B. this will keep executing the aggregation job and the load on the same database instance and this will not resolve loading issues
Not D. using Kinesis Data Firehose to replicate the database is not recommended, the solution should involve DMS. also moving to Kinesis Data Stream for data load requires some changes on the customer side which is not part of the request.

C is the right solution: use DMS to migrate on-premise database, move the aggregation job to the read replica, using Lambda (that supports node.js) behind ALB will not impact client side
upvoted 3 times

shaaam80 1 year, 6 months ago
Selected Answer: C
Answer C
upvoted 1 times

NikkyDicky 1 year, 12 months ago
Selected Answer: C
It's a c
upvoted 1 times

SkyZeroZx 2 years ago
Selected Answer: C
Keyworks = DMS & RDS Proxy
Then C
upvoted 2 times

leehjworking 2 years, 1 month ago
Selected Answer: C
AD: restart = interruption?
B: ASG...Why?
upvoted 3 times

chikorita 2 years ago
why ...oh...why?
upvoted 1 times

mfsec 2 years, 3 months ago
Selected Answer: C
ill go with C
upvoted 1 times

dev112233xx 2 years, 3 months ago
Selected Answer: C
C.. even though question didn't mention the total time of each job. If the job takes more than 15m then Lambda can't be used. Probably the solution with ASG and EC2 is better .. not sure!
upvoted 3 times

zejou1 2 years, 3 months ago
Selected Answer: C
ALB because you are pointing to to Lambda function, not a network address

Look at AWS DMS feature https://aws.amazon.com/dms/features/

Main requirement - needs the migration to occur w/out interruptions or changes to the company's customers.

C keeps it stupid simple w/ no service interruption
upvoted 1 times

vherman 2 years, 3 months ago
Could anybody explain why ALB? I'd go with API Gateway
upvoted 1 times

zejou1 2 years, 3 months ago
Application - you are using Lambda functions that will be sending api commands, you would use network when it is just about routing

upvoted 1 times

Sarutobi 2 years, 4 months ago

Selected Answer: C

I would say C.

upvoted 1 times

hobokabobo 2 years, 4 months ago

I have a feeling that none of the approaches will work.

a) We have two sources that change the database: migration and new data coming in. In a relational database this results in inconsistent data. Constraints will not be fulfilled.

b) until the database is fully synced the second database has inconsistent data. Some parts of relations and parts of entities are still missing. Constraints will not be fulfilled.

None if the approaches addresses that aggregation tasks fail because of inconsistency of the data base.

upvoted 1 times

hobokabobo 2 years, 4 months ago

ACID principle: atomicity, consistency, isolation and durability. All solutions violate this basic principle of relational databases.

https://en.wikipedia.org/wiki/ACID

upvoted 1 times

God_Is_Love 2 years, 4 months ago

Issue could be because of same db used for writing and reading heavily. solution to separate this into

read replica only for reading. DMS for data migration to aws from onpremises.Writing app to DB and Reading app from DB for reports. Writing app needs RDSProxy and saves data.Reading app reads from replica.

B is wrong because, Reading job (aggregation) needs to use replica which is mentioned in C. C is correct.

upvoted 2 times

Fatoch 2 years, 4 months ago

is it C or B?

Same person answers two times two different answers

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: C

C is corect

upvoted 3 times

masetromain 2 years, 5 months ago

Selected Answer: C

C.

This option would meet the requirements of resolving the data loading issue and migrating without interruption or changes for the company's customers. By using AWS DMS for continuous data replication, the company can ensure that the data being migrated is up to date. By setting up an Aurora Replica and moving the aggregation jobs to run against it, the company can offload some of the read workload from the primary database and reduce the risk of issues with the load jobs. By using AWS Lambda functions behind an ALB and Amazon RDS Proxy to write to the Aurora MySQL database, the company can add an extra layer of security and scalability to the data collection process. Finally, by pointing the collector DNS record to the ALB after the databases are synced and disabling the AWS DMS sync task, the company can ensure a smooth cutover to the new environment.

upvoted 4 times

masetromain 2 years, 5 months ago

A.

This option would not work as it would require to change the primary database and also it may cause interruption for the company's customers during the cutover process.

B.

This option would not work as it would not include Aurora Replica to offload the read workload, this would result in aggregation jobs running on the primary database which can cause the load jobs to fail during heavy loads.

D.

This option would not work as it would require to use kinesis data stream which may cause performance issues and also it may not be the best fit for this use case. Additionally, using Kinesis Data Firehose would add complexity to the data replication process, and may result in increased latency or data loss.

upvoted 2 times

Question #51 *Topic 1*

A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled.

Which solution will meet these requirements?

A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key. Use the AWS CLI to re-upload all objects in the S3 bucket. Set an S3 bucket policy to deny unencrypted PutObject requests.

B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests. Use the AWS CLI to re-upload all objects in the S3 bucket.

C. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.

D. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key. Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket. Use the AWS CLI to re-upload all objects in the S3 bucket.

Suggested Answer: *D*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingKMSEncryption.html

So the correct answer is B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests. Use the AWS CLI to re-upload all objects in the S3 bucket.
upvoted 42 times

hamimelon 1 year, 9 months ago

Not B. "must be encrypted by keys that the company's security team manages". This implies the company does not wanna use AWS KMS.
upvoted 5 times

hogtrough 1 year, 4 months ago

This is why they would use Customer-managed keys in AWS KMS. It is absolutely B
upvoted 3 times

jpa8300 1 year, 6 months ago

Hamimmelon, the Company's security Team can manage the AWS KMS service, so B is the right answer. All the others are not valid.
upvoted 2 times

masetromain 2 years, 5 months ago

Option A is not correct because it uses SSE-S3 with a customer-managed key, but it does not specify how the security team will manage the encryption keys. Additionally, it only denies unencrypted PutObject requests but does not specify how the objects will be encrypted.

Option C is not correct because it does not specify how the security team will manage the encryption keys and it does not specify how the objects will be encrypted.

Option D is not correct because it uses AES-256 with a customer-managed key, but it does not specify how the security team will manage the encryption keys. Additionally, it simply denies unencrypted PutObject requests, but it doesn't specify how the objects will be encrypted.
upvoted 8 times

jpa8300 1 year, 6 months ago

And adding to this in option D they specify uses default AES-256, but KMS also uses the same, so this option just don't make sense.
upvoted 1 times

Musk 2 years, 5 months ago

What about the requirement of customer managed keys?

upvoted 10 times

hobokabobo 2 years, 4 months ago

Completely ignores the task to solve: "all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. "

upvoted 4 times

cherep87 2 years, 3 months ago

Use the AWS CLI to re-upload all objects in the S3 bucket. -
https://docs.aws.amazon.com/AmazonS3/latest/userguide/default-bucket-encryption.html
Changes to note before enabling default encryption
After you enable default encryption for a bucket, the following encryption behavior applies:

There is no change to the encryption of the objects that existed in the bucket before default encryption was enabled.

When you upload objects after enabling default encryption:

If your PUT request headers don't include encryption information, Amazon S3 uses the bucket's default encryption settings to encrypt the objects.

upvoted 1 times

hobokabobo 2 years, 2 months ago

Task is to replace any AWS Managed keys to ones "that the company's security team manages"
So they tell us to find a solution that does not use AWS Managed Keys.

upvoted 4 times

hogtrough 1 year, 4 months ago

No, the task was to replace SSE-SE keys which have no relation to AWS KMS.

"Amazon S3 automatically enables server-side encryption with Amazon S3 managed keys (SSE-S3) for new object uploads.

Unless you specify otherwise, buckets use SSE-S3 by default to encrypt objects. However, you can choose to configure buckets to use server-side encryption with AWS Key Management Service (AWS KMS) keys (SSE-KMS) instead. "

https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingKMSEncryption.html

upvoted 1 times

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: D

I think D is correct.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerSideEncryptionCustomerKeys.html

upvoted 21 times

djeong95 1 year, 4 months ago

The issue with D is that it doesn't make it clear where the encryption is happening like all the other options do. Is it server-side (we assume that it is, but it is not what is written)? Or is it client-side?

upvoted 1 times

Curious76 Most Recent 1 month, 1 week ago

Selected Answer: A

A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key. Use the AWS CLI to re-upload all objects in the S3 bucket. Set an S3 bucket policy to deny unencrypted PutObject requests.

Here's why:
Requirements Recap:
All current and future objects must be encrypted using customer-managed keys.

The current encryption is SSE-S3, which uses S3-managed keys, not customer-managed keys.

The bucket does not have versioning enabled, so overwriting (re-uploading) is necessary to change encryption on existing objects.

upvoted 1 times

amministrazione 10 months ago

B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests. Use the AWS CLI to re-upload all objects in the S3 bucket.

upvoted 1 times

Jason666888 11 months ago

Selected Answer: B

AWS KMS (Key Management Service) allows for customer-managed keys (CMKs), which can indeed be considered as "keys that the company's security team manages"

upvoted 1 times

Helpnosense 1 year ago

Selected Answer: B

In s3 option there is no option to select AES256 custom key.

upvoted 1 times

higashikumi 1 year, 1 month ago

Selected Answer: B

To meet the requirement for encrypting all current and future objects in an Amazon S3 bucket with keys managed by the company's security team, change the S3 bucket's default encryption to server-side encryption with AWS KMS managed keys (SSE-KMS). Implement an S3 bucket policy to deny unencrypted PutObject requests, ensuring all new uploads are encrypted with the specified KMS key. Then, use the AWS CLI to re-upload all existing objects to the S3 bucket, enforcing the new encryption policy on current data. This approach ensures compliance by applying KMS encryption to both new and existing objects without causing disruptions .

upvoted 1 times

Malcnorth59 1 year, 1 month ago

Selected Answer: B

The solutions need to use SSE-KMS so that the security team can manage the keys, but they also need to ensure that current and future objects are encrypted using customer-managed keys.

upvoted 1 times

TonytheTiger 1 year, 1 month ago

Selected Answer: D

Not Option D: " Amazon S3 server-side encryption uses 256-bit Advanced Encryption Standard Galois/Counter Mode (AES-GCM) to encrypt all uploaded objects." AES- 256 is already the default, so you can't change it.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingServerSideEncryption.html

upvoted 1 times

mav3r1ck 1 year, 3 months ago

Selected Answer: B

Correct Approach: This option is accurate and meets all the specified requirements. By changing the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS), the company can use customer managed keys (CMKs) for encryption. This allows the security team to manage the keys, addressing the core requirement.
Setting an S3 bucket policy to deny unencrypted PutObject requests ensures future compliance with the encryption policy.
Re-uploading all objects using the AWS CLI ensures that existing objects are encrypted under the new policy, making sure that both current and future objects are encrypted with the keys managed by the company's security team.

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B, option D confuses encryption options. AES-256 is part of the SSE-S3 encryption method and doesn't directly involve customer-managed keys

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: B

The solution that meets the requirements for encrypting all current and future objects in the Amazon S3 bucket with keys that the company's security team manages, while ensuring server-side encryption, is:
Option B is correct because it directly addresses the new requirement by changing the default encryption method to SSE-KMS, which allows the use of AWS Key Management Service (KMS) keys managed by the company's security team. This option ensures that all future uploads are encrypted with the specified KMS key. It also includes re-uploading existing objects to ensure they are encrypted under the new scheme. Setting an S3 bucket policy to deny unencrypted PutObject requests enforces the encryption requirement for all new uploads.

upvoted 1 times

8608f25 1 year, 4 months ago

Option D is incorrect because it refers to “AES-256 with a customer managed key” in a way that mixes concepts. AES-256 is the encryption standard used by SSE-S3 and does not directly apply to the use of customer managed keys. For managing keys, the correct approach is through SSE-KMS, which allows specifying a customer managed AWS KMS key.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

Not A. SSE-S3 with a customer managed key is not an actual option as SSE-S3 uses S3 managed keys
Not C. S3 bucket policy cannot automatically encrypt objects on GetObject and PutObject requests. With policies you can only allow/deny actions from specific principals
Not D. AES-256 with a customer managed key is not an actual option as AES-256 is used as value for the header x-amz-server-side-encryption to set SSE-S3 on putObject and SSE-S3 uses S3 managed keys

B is correct as server-side encryption with AWS KMS managed encryption keys (SSE-KMS) is an actual default encryption settings for S3 bucket and you can use S3 bucket policy to deny unencrypted PutObject. These ensure all new objects will be encrypted with customer managed keys. Then using aws cli to re-upload all object will overwrite existing objects (versioning is not enabled)

upvoted 2 times

ismeagain 1 year, 6 months ago

Selected Answer: D

i think D is correct as B is mentioned KMS managed key..

upvoted 1 times

Impromptu 1 year, 6 months ago

Selected Answer: B

A - You cannot define your own key
B - Correct. Using SSE-KMS and your own KMS customer managed key, you adhere to the requirements
C - Does not encrypt existing objects, and you cannot "change" the request to "automatically" encrypt
D - You can only choose between SSE-S3 and SSE-KMS (or now DSSE-KMS as well) for default encryption. Underlying the SSE-S3 refers to AES-256 (cfr. "s3:x-amz-server-side-encryption": "AES256") but you cannot specify your customer managed key in that case.

upvoted 1 times

_Juwon 1 year, 6 months ago

Selected Answer: B

If use KMS-CMK , wouldn't it be possible to manage keys directly while using KMS? Does anyone have an opinion on this?

upvoted 1 times

eurriola10 1 year, 7 months ago

Selected Answer: B

B is correct
https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingKMSEncryption.html#aws-managed-customer-managed-keys

When you use server-side encryption with AWS KMS (SSE-KMS), you can use the default AWS managed key, or you can specify a customer managed key that you have already created.

upvoted 2 times

Question #52 *Topic 1*

A company is running a web application in the AWS Cloud. The application consists of dynamic content that is created on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group that is configured as a target group for an Application Load Balancer (ALB).

The company is using an Amazon CloudFront distribution to distribute the application globally. The CloudFront distribution uses the ALB as an origin. The company uses Amazon Route 53 for DNS and has created an A record of www.example.com for the CloudFront distribution.

A solutions architect must configure the application so that itis highly available and fault tolerant.

Which solution meets these requirements?

A. Provision a full, secondary application deployment in a different AWS Region. Update the Route 53 A record to be a failover record. Add both of the CloudFront distributions as values. Create Route 53 health checks.

B. Provision an ALB, an Auto Scaling group, and EC2 instances in a different AWS Region. Update the CloudFront distribution, and create a second origin for the new ALCreate an origin group for the two origins. Configure one origin as primary and one origin as secondary.

C. Provision an Auto Scaling group and EC2 instances in a different AWS Region. Create a second target for the new Auto Scaling group in the ALB. Set up the failover routing algorithm on the ALB.

D. Provision a full, secondary application deployment in a different AWS Region. Create a second CloudFront distribution, and add the new application setup as an origin. Create an AWS Global Accelerator accelerator. Add both of the CloudFront distributions as endpoints.

Suggested Answer: *B*

Community vote distribution

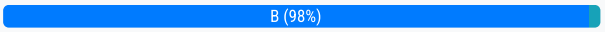

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct answer is B. Provisioning an ALB, an Auto Scaling group, and EC2 instances in a different AWS region provides redundancy and failover capability for the application. By creating a second origin for the new ALB in the second region, the CloudFront distribution can automatically route traffic to the healthy origin in case of an issue with the primary origin. This ensures that the application remains highly available and fault-tolerant.

Option A is not correct because it uses Route 53 failover records, which can result in increased latency and DNS resolution time for clients. Option C is not correct because it doesn't provide redundancy for the load balancer, which is a critical component of the application. Option D is not correct because it does not provide redundancy for the application in case of an issue with the primary origin in the first region.

upvoted 27 times

God_Is_Love Highly Voted 2 years, 4 months ago

For HA, always user second region but its there in all options. Here Cloudfront distribution multiple origin groups is the key point Solution Architects should know of. Configuring 2nd origin as ALB --> EC2 instances target group in another regions setup makes highly available. If Cloudfront detects that response is Http error (fault) code like 4XX,5XX etc, it will failover to secondary origin (ALB of another region) which makes this fault tolerant. Answer is B.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

upvoted 11 times

amministrazione Most Recent 10 months ago

B. Provision an ALB, an Auto Scaling group, and EC2 instances in a different AWS Region. Update the CloudFront distribution, and create a second origin for the new ALCreate an origin group for the two origins. Configure one origin as primary and one origin as secondary.

upvoted 1 times

8693a49 11 months ago

Selected Answer: A

This architecture is an active-active DR strategy. You would do it with R53 failover because R53 has healthchecks, and once the primary is down all requests go to the failover. With CloudFront failover, all requests would continue to hit the failed primary before being routed to the failover distribution, which increases latency and possibly compounds problems in the failed stack. Interestigly, the best solution would actually be a combination between A and B, as this blog post shows:

https://aws.amazon.com/blogs/networking-and-content-delivery/improve-web-application-availability-with-cloudfront-and-route53-hybrid-origin-failover/

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: B

A is wrong because CloudFront distros can't be added to Route 53.
B is correct
C is wrong because ALBs are single region and don't do failover.
D would work, but is overengineered in this context.

upvoted 2 times

8693a49 11 months ago

You can add CloudFront distros to R53 using alias records: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfront-distribution.html

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: B

Option B is correct because it involves creating a redundant setup in another AWS Region with its own ALB, Auto Scaling group, and EC2 instances. By updating the CloudFront distribution to include a second origin for the new ALB and creating an origin group with primary and secondary origins, CloudFront can automatically route traffic to the secondary origin if the primary is unhealthy. This setup leverages CloudFront's global reach to improve availability and fault tolerance without the need for DNS-level changes.
Option A is not correct because it suggests creating a secondary deployment and updating the Route 53 A record to be a failover record with both CloudFront distributions as values. While Route 53 health checks and failover records can improve availability, CloudFront distributions themselves cannot be directly specified as values in A records for failover purposes. This option might lead to confusion in its implementation details.

upvoted 2 times

bjexamprep 1 year, 5 months ago

Selected Answer: B

Who the hell cooked this terrible question design.
Usually, HA means single region, DR means cross region. The question is asking HA while all the answer are using cross region solutions.
When Dynamic content is involved, the dynamic content has to be store in a persistent storage, while question says the dynamic content is store on the EC2 instances in an ASG, which means the EC2 instances are ephemeral.
And when Dynamic content is involved, no matter HA or DR, a replication component must be built so that the Dynamic content will be replicated to the other side so that it can be available when the event happens. While, none of the answers mentions replication at all.

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

Not A. CloudFront is a global service, having two distributions will not increase fault-tolerance
Not C. Single ALB is a single-point-of-failure and also you cannot have Target Group in a different region
Not D. CloudFront is a global service, having two distributions will not increase fault-tolerance and combining CloudFront with AWS Global Accelerator makes no sense

B is correct as provisioning an ALB, an Auto Scaling group, and EC2 instances in a different AWS region provides redundancy and failover capability for the application. The origin group is the right way to enable failover for CloudFront distributions origin

upvoted 3 times

holymancolin 1 year, 7 months ago

Selected Answer: B

Not Create a second CloudFront Distribution, it's update the distribution with multi origins.

Ref:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html#concept_origin_groups.creating
"Make sure the distribution has more than one origin. If it doesn't, add a second origin."

upvoted 1 times

NikkyDicky 1 year, 12 months ago

it's a B

upvoted 1 times

[Removed] 2 years ago

Selected Answer: B

Both A and B would work, but A is tangibly worse in terms of performing fail-over (because it relies on DNS) and gains you little, since CloudFront is highly available by its nature, making a second CF distribution doesn't improve your application's robustness.

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: B

Provision an ALB, an Auto Scaling group, and EC2 instances in a different AWS Region.

upvoted 1 times

dev112233xx 2 years, 3 months ago

Selected Answer: B

B is the best solution with very high availability (compared to the R53 failover solution)

upvoted 1 times

Ajani 2 years, 3 months ago

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

upvoted 1 times

Sarutobi 2 years, 4 months ago

Selected Answer: B

B looks good.

upvoted 1 times

masssa 2 years, 5 months ago

Selected Answer: B

B is correct.

C is not correct, because ALB is regional service, so ALB have to be added too.

upvoted 2 times

Question #53 *Topic 1*

A company has an organization in AWS Organizations that has a large number of AWS accounts. One of the AWS accounts is designated as a transit account and has a transit gateway that is shared with all of the other AWS accounts. AWS Site-to-Site VPN connections are configured between all of the company's global offices and the transit account. The company has AWS Config enabled on all of its accounts.

The company's networking team needs to centrally manage a list of internal IP address ranges that belong to the global offices. Developers will reference this list to gain access to their applications securely.

Which solution meets these requirements with the LEAST amount of operational overhead?

A. Create a JSON file that is hosted in Amazon S3 and that lists all of the internal IP address ranges. Configure an Amazon Simple Notification Service (Amazon SNS) topic in each of the accounts that can be invoked when the JSON file is updated. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with the updated IP address ranges.

B. Create a new AWS Config managed rule that contains all of the internal IP address ranges. Use the rule to check the security groups in each of the accounts to ensure compliance with the list of IP address ranges. Configure the rule to automatically remediate any noncompliant security group that is detected.

C. In the transit account, create a VPC prefix list with all of the internal IP address ranges. Use AWS Resource Access Manager to share the prefix list with all of the other accounts. Use the shared prefix list to configure security group rules in the other accounts.

D. In the transit account, create a security group with all of the internal IP address ranges. Configure the security groups in the other accounts to reference the transit account's security group by using a nested security group reference of " /sg-1a2b3c4d".

Suggested Answer: *C*

Community vote distribution

C (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct answer is option C. In this solution, a VPC prefix list is created in the transit account with all of the internal IP address ranges, and then shared to all of the other accounts using AWS Resource Access Manager. This allows for central management of the IP address ranges, and eliminates the need for manual updates to security group rules in each account. This solution also allows for compliance checks to be run using AWS Config and for any non-compliant security groups to be automatically remediated.

Option A is not correct because it would require manual updates to the JSON file and would also require developers to manually update their security group rules, which would lead to operational overhead.

Option B is not correct because it would require the creation of a new AWS Config managed rule and it would also require manual updates to the security group rules in each account.

Option D is not correct because it would require manual updates to the security group in the transit account and it would also lead to operational overhead.

upvoted 24 times

jpa8300 1 year, 6 months ago

I agree that option C is probable the best one, but B is also correct, there is no manual updates to the SG, the remediation is automated in ASW Config. In option C you also need to manual update the prefix list, no? Imagine a new CIDR appears in the offices.

upvoted 1 times

chicagobeef 1 year, 5 months ago

I doubt all the security groups in the accounts will use the same CIDR ranges. They just need a way to centrally manage the CIDR prefixes. The question did not say that everyone has to comply and any non-compliant resources needs to be remdiated.

upvoted 2 times

Aritra88 Most Recent 6 months, 3 weeks ago

Selected Answer: C

A VPC Prefix List is a reusable, user-defined resource in Amazon Virtual Private Cloud (VPC) that contains a collection of IP address ranges. These ranges can represent destinations or sources for traffic, and the prefix list can be referenced in various configurations like security groups, route

tables, or network ACLs.
upvoted 1 times

Tiger4Code 7 months ago

Selected Answer: C

C: in the shared account create a VPC Prefix list, share it using RAM, then SGs can reference it
upvoted 1 times

amministrazione 10 months ago

C. In the transit account, create a VPC prefix list with all of the internal IP address ranges. Use AWS Resource Access Manager to share the prefix list with all of the other accounts. Use the shared prefix list to configure security group rules in the other accounts.
upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

Not A. This requires to maintain the JSON file, SNS topic in each account, Lambda to update SG. This is a lot of work, also not clear what accounts holds the S3 with the JSON
Not B. I was not able to spot a managed AWS Config rule that could help in this case https://docs.aws.amazon.com/config/latest/developerguide/managed-rules-by-aws-config.html (but I do not recall managed rule by hart and this doesn't sound like a remote use case, so in the exam this could trick me)
upvoted 2 times

ninomfr64 1 year, 5 months ago

Not D. You can reference a VPC SG in other account VPCs when you have VPC peering in place, this is not mentioned in the scenario https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-security-groups.html. Since there is a Transit Gateway involved it is unlikely to have VPC peering and the resources in a VPC attached to a transit gateway cannot access the security groups of a different VPC that is also attached to the same transit gateway https://docs.aws.amazon.com/vpc/latest/tgw/tgw-vpc-attachments.html (this option initially was not bad for me)

C works well as prefix lists are created exactly for this purpose https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html
upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C for sure
upvoted 1 times

Asds 2 years ago

Selected Answer: C

Definitely prefix
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

prefix list and RAM
upvoted 2 times

dev112233xx 2 years, 3 months ago

Selected Answer: C

C makes sense ✅
upvoted 2 times

zozza2023 2 years, 5 months ago

Selected Answer: C

https://www.examtopics.com/discussions/amazon/view/82131-exam-aws-certified-solutions-architect-professional-topic-1/
upvoted 2 times

AjayD123 2 years, 5 months ago

Selected Answer: C

https://aws.amazon.com/blogs/networking-and-content-delivery/simplify-network-routing-and-security-administration-with-vpc-prefix-lists/#:~:text=A%20Prefix%20List%20is%20a,Resource%20Access%20Manager%20(RAM).
upvoted 4 times

Question #54 *Topic 1*

A company runs a new application as a static website in Amazon S3. The company has deployed the application to a production AWS account and uses Amazon CloudFront to deliver the website. The website calls an Amazon API Gateway REST API. An AWS Lambda function backs each API method.

The company wants to create a CSV report every 2 weeks to show each API Lambda function's recommended configured memory, recommended cost, and the price difference between current configurations and the recommendations. The company will store the reports in an S3 bucket.

Which solution will meet these requirements with the LEAST development time?

A. Create a Lambda function that extracts metrics data for each API Lambda function from Amazon CloudWatch Logs for the 2-week period. Collate the data into tabular format. Store the data as a .csv file in an S3 bucket. Create an Amazon EventBridge rule to schedule the Lambda function to run every 2 weeks.

B. Opt in to AWS Compute Optimizer. Create a Lambda function that calls the ExportLambdaFunctionRecommendations operation. Export the .csv file to an S3 bucket. Create an Amazon EventBridge rule to schedule the Lambda function to run every 2 weeks.

C. Opt in to AWS Compute Optimizer. Set up enhanced infrastructure metrics. Within the Compute Optimizer console, schedule a job to export the Lambda recommendations to a .csv file. Store the file in an S3 bucket every 2 weeks.

D. Purchase the AWS Business Support plan for the production account. Opt in to AWS Compute Optimizer for AWS Trusted Advisor checks. In the Trusted Advisor console, schedule a job to export the cost optimization checks to a .csv file. Store the file in an S3 bucket every 2 weeks.

Suggested Answer: *B*

Community vote distribution

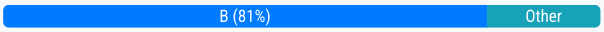

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct answer is B. Opting in to AWS Compute Optimizer and creating a Lambda function that calls the ExportLambdaFunctionRecommendations operation is the least development time solution. This option allows you to use the built-in AWS Compute Optimizer service to extract metrics data and export it as a CSV file, which can then be stored in an S3 bucket.

Option A is not correct because it requires the development of a Lambda function that extracts metrics data and collates it into tabular format, which adds development time. Option C is not correct because it requires the setup of enhanced infrastructure metrics, which adds development time. Option D is not correct because it requires purchasing the AWS Business Support plan and using the Trusted Advisor console, which adds development time.

upvoted 23 times

zozza2023 Highly Voted 2 years, 5 months ago

Selected Answer: B

AWS compute optimizer+ lambda

upvoted 9 times

Aritra88 Most Recent 6 months, 3 weeks ago

Selected Answer: B

Answer B

Solution Steps

1. Use AWS Compute Optimizer for Lambda Recommendations

AWS Compute Optimizer provides recommendations for Lambda functions, including:

* Recommended memory size to improve performance or reduce cost.

* Current and recommended cost comparisons.

You can query AWS Compute Optimizer using the AWS Management Console, AWS CLI, or SDKs to retrieve the necessary data for your report.

2. Automate Data Retrieval

Set up an AWS Lambda function to automate the process:

1. Query Compute Optimizer:

* Use the GetLambdaFunctionRecommendations API to retrieve:

* Current memory size
* Recommended memory size
* Current and recommended cost

upvoted 1 times

amministrazione 10 months ago

B. Opt in to AWS Compute Optimizer. Create a Lambda function that calls the ExportLambdaFunctionRecommendations operation. Export the .csv file to an S3 bucket. Create an Amazon EventBridge rule to schedule the Lambda function to run every 2 weeks.

upvoted 1 times

8693a49 11 months ago

Why would anyone need to memorize whether Compute Optimizer reports can be scheduled from the UI or must be done through API calls? This is so unnecessary *rolls eyes

upvoted 3 times

khchan123 1 year, 3 months ago

Selected Answer: C

The correct answer is C.
Option A involves creating a custom Lambda function to extract metrics data from CloudWatch Logs and generate the CSV report, which would require more development time compared to using the Compute Optimizer service.

Option B is partially correct, as it involves using Compute Optimizer and a Lambda function, but it misses the ability to schedule recurring exports directly within the Compute Optimizer console.

Option D suggests using AWS Trusted Advisor, which is a service for monitoring best practices and resources, but it does not provide the specific Lambda function memory and cost recommendations required in this scenario.

upvoted 3 times

helloworldabc 10 months ago

just B

upvoted 2 times

8608f25 1 year, 4 months ago

Selected Answer: B

Option B is the most efficient and straightforward solution. By opting into AWS Compute Optimizer, the company can leverage AWS's service for recommendations on optimal AWS resource configurations based on utilization metrics. Using the ExportLambdaFunctionRecommendations operation allows for automating the retrieval of the desired optimization data with minimal code. Scheduling this operation with an Amazon EventBridge rule to run every 2 weeks and exporting the results directly to a CSV file in an S3 bucket meets all the stated requirements with minimal development effort.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

Not A. This requires some serious development, also not 100% sure CW Logs alone provides all the required info.
Not B. This requires some coding to to call the ExportLambdaFunctionRecommendations API
Not D. To create CSV reports (organizational view reports) in Trusted Advisor you need to enable Trusted Advisor in you organization, and AWS Organization is not mentioned in the scenario https://docs.aws.amazon.com/awssupport/latest/user/organizational-view.html

C is the right solution as it allows to schedule report with the required info with no development https://docs.aws.amazon.com/compute-optimizer/latest/ug/exporting-recommendations.html. This is was misleading for me as it mentions to set up enhanced infrastructure metrics that is only available for EC2, but you can do it without development (you can do it from console), this add cost but the ask focus on development effort.

upvoted 2 times

8608f25 1 year, 4 months ago

It is not C. Option C describes using AWS Compute Optimizer and setting up a job within the Compute Optimizer console. However, as of the last update, Compute Optimizer does not provide a direct scheduling feature within the console for exporting recommendations to a CSV file. This option suggests functionality that is not directly available in Compute Optimizer.

upvoted 5 times

AWSCertification2024 1 year, 5 months ago

Selected Answer: B

B is correct
Not C because Enhanced infrastructure metrics is a paid feature of Compute Optimizer that applies to Amazon EC2 instances and instances that are

part of Auto Scaling groups.

upvoted 3 times

enk 1 year, 7 months ago

Selected Answer: D

Lambda = development. Option D has no development. If you are not familiar with dev'ing - publishing a simple Lambda function can require you to wrap all the Node.js or Python or whatever programming language libraries with it in order to execute correctly within AWS Lambda. Configuring Trusted Advisor (GUI) or scheduling a job is NOT considered Development.

upvoted 3 times

KCjoe 1 year, 8 months ago

Selected Answer: D

Basic plan of Trusted Advisor only has 7 core checks. Business plan has all these, so with LEAST development, it must be business plan.
Check categories
Cost optimization
Performance
Security
Fault tolerance
Service limits

upvoted 5 times

rlf 1 year, 8 months ago

B.
Option C is not correct because "Enhanced infrastructure metrics is a paid feature of Compute Optimizer that applies to Amazon EC2 instances."
https://docs.aws.amazon.com/compute-optimizer/latest/ug/enhanced-infrastructure-metrics.html

upvoted 1 times

awsent 1 year, 9 months ago

Selected Answer: B

Computer Optimizer could generate Export for Lambda Functions one-time. In order to schedule every 2 weeks, EventBridge Scheduler/Schedule Rule should be used.

upvoted 4 times

awsent 1 year, 9 months ago

Answer: B
https://aws.amazon.com/blogs/compute/optimizing-aws-lambda-cost-and-performance-using-aws-compute-optimizer/

upvoted 1 times

Simon523 1 year, 9 months ago

Selected Answer: B

AWS Compute Optimizer helps avoid overprovisioning and underprovisioning four types of AWS resources—Amazon Elastic Compute Cloud (EC2) instance types, Amazon Elastic Block Store (EBS) volumes, Amazon Elastic Container Service (ECS) services on AWS Fargate, and AWS Lambda functions—based on your utilization data.

upvoted 4 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

its a B

upvoted 1 times

EricZhang 2 years, 1 month ago

B - https://docs.aws.amazon.com/compute-optimizer/latest/APIReference/API_ExportLambdaFunctionRecommendations.html

upvoted 3 times

Question #55 *Topic 1*

A company's factory and automation applications are running in a single VPC. More than 20 applications run on a combination of Amazon EC2, Amazon Elastic Container Service (Amazon ECS), and Amazon RDS.

The company has software engineers spread across three teams. One of the three teams owns each application, and each time is responsible for the cost and performance of all of its applications. Team resources have tags that represent their application and team. The teams use IAM access for daily activities.

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. A solutions architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.

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

A. Activate the user-define cost allocation tags that represent the application and the team.

B. Activate the AWS generated cost allocation tags that represent the application and the team.

C. Create a cost category for each application in Billing and Cost Management.

D. Activate IAM access to Billing and Cost Management.

E. Create a cost budget.

F. Enable Cost Explorer.

Suggested Answer: *ACF*

Community vote distribution

ACF (56%) ADF (44%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: ACF

A, C and F are the correct answers because they provide the required cost reports and analysis for the company's applications and teams.

A. Activating user-defined cost allocation tags that represent the application and the team allows the company to assign costs to specific applications and teams. This allows the company to see how much each application and team is costing them, which is important for cost forecasting and budgeting.

C. Creating a cost category for each application in Billing and Cost Management allows the company to group costs by application. This makes it easier to understand the costs associated with each application and to compare the costs of different applications over time.

F. Enabling Cost Explorer allows the company to analyze costs and usage over time, and to create custom reports and forecasts. This is important for understanding the costs associated with each application and team, and for forecasting future costs.

upvoted 42 times

masetromain 2 years, 5 months ago

B is not correct because AWS generated cost allocation tags are automatically created for some AWS resources, but it does not provide the required cost reports and analysis for the company's applications and teams.

Option D is not correct because IAM access controls are used to limit access to the billing and cost management features, but it is not necessary to configure it to meet the requirements.

E is not correct because Creating a cost budget allows the company to set a budget for their costs and to receive alerts when costs exceed the budget, but it does not provide the required cost reports and analysis for the company's applications and teams.

upvoted 7 times

a_c_ 2 years, 1 month ago

With out granting IAm Access, IAM users cannot access Billing console, so s cannot see the Cost explorer
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/control-access-billing.html.

Question says teams are responsible for cost
I
upvoted 10 times

djeong95 1 year, 4 months ago
In addition to the IAM access problem answer ACF will face, the problem statement already presents us with the information that resources are already tagged by team/application. Creating cost category seems redundant and even if you did create this redundancy, you are faced with the IAM access problem.

If each team is responsible for the cost and the performance, they would need access to the billing console for their team.
upvoted 2 times

e4bc18e 1 year, 2 months ago
So you are wrong, tags can be applied to applications so you can easily find them but unless they are actually activated as user defined billing tags then you will not be able to use those tags in cost analysis. Also you have to enable cost explorer it is not enabled by default and cost explorer lets you see the previous 12 months and creates projections for the next 12, so without that option you will not meet the objective.
upvoted 3 times

spd Highly Voted 2 years, 4 months ago
Selected Answer: ADF
Correct ADF - SInce resources are tagged, C may not require ?
upvoted 18 times

eesa Most Recent 2 months, 1 week ago
Selected Answer: ACF
✓ A. Activate the user-defined cost allocation tags

User-defined tags must be explicitly activated in AWS Billing and Cost Management for cost allocation.

Since the teams already have tags on their resources representing their applications and teams, activating these user-defined tags allows AWS to organize the costs accordingly.

✓ C. Create a cost category for each application

Cost categories simplify grouping related costs.

You can group resources by application or team, making it easier to manage and report on costs for comparison and forecasting.

✓ F. Enable Cost Explorer

Cost Explorer provides historical cost reports and visualizes trends over the past 12 months.

It also allows cost forecasting for the next 12 months, directly fulfilling the requirement to compare and forecast costs.
upvoted 1 times

bhanus 7 months ago
Selected Answer: ACF
ACF
User defined tags to separate out billing.
Grouping the costs.
Cost Explorer to analyze the costs.
upvoted 1 times

amministrazione 10 months ago
A. Activate the user-define cost allocation tags that represent the application and the team.
C. Create a cost category for each application in Billing and Cost Management.
E. Create a cost budget
upvoted 3 times

neta1o 11 months, 2 weeks ago
Selected Answer: ACF

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. - Tagging and Cost Categories

The company also must be able to create reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. -Cost Explorer

upvoted 3 times

alex_heavy 12 months ago

Selected Answer: ADF

A. User defined cost allocation tags: application, team

D. Activate IAM access to Billing and Cost Management:

"The teams use IAM access for daily activities."

https://docs.aws.amazon.com/cost-management/latest/userguide/control-access-billing.html

F. Enable Cost Explorer: https://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html

C is NOT needed, because A already will give a usage view by "tags that represent their application and team" "The company needs to determine which costs on the monthly AWS bill are attributable to each application or team"

upvoted 4 times

TonytheTiger 1 year, 2 months ago

Selected Answer: ACF

Option ACF and NOT ADF - Cost allocation helps you identify who is spending what, within your organization. Cost categories is a cost allocation service to help you map your AWS costs, to your unique internal business structures.

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html

upvoted 3 times

mav3r1ck 1 year, 3 months ago

Selected Answer: ACF

Focusing on enabling the company to attribute AWS costs to each application or team, create cost comparison reports for the last 12 months, and forecast costs for the next 12 months,..Answer: A, C, F.

upvoted 2 times

mav3r1ck 1 year, 3 months ago

Explanation of Exclusions: B, D, F

upvoted 1 times

mav3r1ck 1 year, 3 months ago

E. Create a cost budget: Creating a cost budget is valuable for managing expenses and avoiding overspending, but it does not directly facilitate the attribution of costs to applications or teams, nor does it aid in the creation of historical comparison reports or forecasts in the manner required by the company.

upvoted 2 times

mav3r1ck 1 year, 3 months ago

D. Activate IAM access to Billing and Cost Management: While important for ensuring that team members can access billing information, this action itself doesn't contribute directly to organizing or reporting on costs by application or team, nor does it facilitate forecasting.

upvoted 1 times

mav3r1ck 1 year, 3 months ago

[correction for typo error above] Explanation of Exclusions: B, D, E

upvoted 1 times

mav3r1ck 1 year, 3 months ago

here's the detailed recommendation:

upvoted 1 times

mav3r1ck 1 year, 3 months ago

A. Activate user-defined cost allocation tags: User-defined tags need to be activated for cost allocation purposes. These tags, representing applications and teams, are crucial for attributing costs accurately to the responsible entities within the company. Once activated, these tags will appear in the AWS Billing and Cost Management dashboard, enabling detailed tracking and reporting based on the specified tags.

upvoted 1 times

mav3r1ck 1 year, 3 months ago

F. Enable Cost Explorer: Cost Explorer is essential for analyzing past spending and forecasting future costs. It allows for detailed reports that can compare costs from the last 12 months and helps in forecasting for the next 12 months. With the data segmented by user-defined cost allocation tags, Cost Explorer can provide the insights needed to meet the company's reporting and forecasting requirements.
C. Create a cost category for each application in Billing and Cost Management: Cost categories allow for the organization of cost and usage data into logical groups that reflect the company's internal structure, such as by application or team. By leveraging the user-defined tags activated in step A, cost categories can automate the process of cost attribution to these entities, simplifying the creation of targeted reports and forecasts.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: ACF

Agree with ACF

upvoted 3 times

Dgix 1 year, 3 months ago

Selected Answer: ACF

For the full granularity, C is needed rather than D.

upvoted 3 times

a54b16f 1 year, 3 months ago

Selected Answer: ADF

C is not needed. Option A activated the tag, so we could use tags to generate reports. There is no need to create cost category for individual applications, which could be a huge effort and not practical, what if you have hundreds of applications...

upvoted 3 times

a54b16f 1 year, 4 months ago

Selected Answer: ADF

Correct ADF - SInce resources are tagged

upvoted 2 times

8608f25 1 year, 4 months ago

Selected Answer: ACF

Correct answers are:
A. Activate the user-defined cost allocation tags that represent the application and the team. User-defined cost allocation tags allow you to organize your AWS bill by categorizing costs according to your business's organizational structures (e.g., by application or team).
C. Create a cost category for each application in Billing and Cost Management. Cost categories enable you to create custom groupings of your AWS costs. By creating a cost category for each application, you can group costs more granularly, which is helpful for detailed reporting and cost attribution to specific teams or applications.
F. Enable Cost Explorer. Cost Explorer is a tool that allows you to visualize, understand, and manage your AWS costs and usage over time. By enabling Cost Explorer, you can create detailed reports to compare costs from the last 12 months and forecast costs for the next 12 months, meeting the company's requirements for cost management and planning.

upvoted 2 times

8608f25 1 year, 4 months ago

Option B is not correct. It refers to activating AWS generated cost allocation tags. While AWS-generated tags can provide useful information, they do not typically represent specific applications or teams unless those entities are directly associated with AWS-defined resources or actions. For custom application and team tracking, user-defined tags (Option A) are more appropriate.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: ADF

Not B. AWS generated tags do not allow you to identify app. You need user-defined tags for this
Not C. Cost Categories allows to define rule to group costs into categories using different dimensions such as: account, tag, service, charge type, and other cost categories. In this scenario User-defined tags are enough to identify applications and teams.
Not E. Budget doesn't help you in creating reports to compare costs from the last 12 months and to help forecast costs for the next 12 months. Use Cost Explorer instead-

upvoted 5 times

jpa8300 1 year, 5 months ago

Selected Answer: ADF

See below severlight explanation. I agree with it.

upvoted 3 times

Dips3009 1 year, 6 months ago

can someone help me with this solutions, as I am confused between ACF and ADF

upvoted 1 times

Question #56 *Topic 1*

An AWS customer has a web application that runs on premises. The web application fetches data from a third-party API that is behind a firewall. The third party accepts only one public CIDR block in each client's allow list.

The customer wants to migrate their web application to the AWS Cloud. The application will be hosted on a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in a VPC. The ALB is located in public subnets. The EC2 instances are located in private subnets. NAT gateways provide internet access to the private subnets.

How should a solutions architect ensure that the web application can continue to call the third-party API after the migration?

A. Associate a block of customer-owned public IP addresses to the VPC. Enable public IP addressing for public subnets in the VPC.

B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC.

C. Create Elastic IP addresses from the block of customer-owned IP addresses. Assign the static Elastic IP addresses to the ALB.

D. Register a block of customer-owned public IP addresses in the AWS account. Set up AWS Global Accelerator to use Elastic IP addresses from the address block. Set the ALB as the accelerator endpoint.

Suggested Answer: *B*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct solution is B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC. This will ensure that the web application can continue to call the third-party API after the migration by using the customer-owned public IP addresses that were assigned to the NAT gateways. This ensures that the third-party API will only see traffic coming from the customer-owned IP addresses that are on the allow list. Option A,C and D doesn't make sense in this context.

upvoted 20 times

amministrazione Most Recent 10 months ago

B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

In this scenario EC2 instances access the 3P APIs via NAT Gateway. 3P API FW see IP of the NAT Gateway. You can assign Elastic IP to NAT Gateway and you can allocate an IP address from a pool that you have brought to your AWS account to the Elastic IP. Thus B is correct.
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

its a B

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

KEYWORD = NAT gateways in the VPC

upvoted 2 times

AWS_Sam 2 years, 1 month ago

B is the only option that makes sense.

upvoted 1 times

SkyZeroZx 2 years, 1 month ago

Selected Answer: B

B make sense

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

Register a block of customer owned public IP's

upvoted 2 times

dev112233xx 2 years, 3 months ago

Selected Answer: B

B is the only solution

upvoted 2 times

zozza2023 2 years, 5 months ago

Selected Answer: B

The correct solution is B

upvoted 4 times

Question #57 *Topic 1*

A company with several AWS accounts is using AWS Organizations and service control policies (SCPs). An administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAllActions",
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*"
        },
        {
            "Sid": "DenyCloudTrail",
            "Effect": "Deny",
            "Action": "cloudtrail:*",
            "Resource": "*"
        }
    ]
}
```

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the administrator address this problem?

A. Add s3:CreateBucket with "Allow" effect to the SCP.

B. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.

C. Instruct the developers to add Amazon S3 permissions to their IAM entities.

D. Remove the SCP from account 1111-1111-1111.

Suggested Answer: *C*

Community vote distribution

C (89%) 11%

Atila50 Highly Voted 2 years, 5 months ago

Selected Answer: C

SCP doesn't grant permission

upvoted 23 times

c73bf38 2 years, 4 months ago

Per the DOCS:

Service control policies (SCPs) are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization. SCPs help you to ensure your accounts stay within your organization's access control guidelines. SCPs are available only in an organization that has all features enabled. SCPs aren't available if your organization has enabled only the consolidated billing features. For instructions on enabling SCPs, see Enabling and disabling policy types.

upvoted 7 times

c73bf38 2 years, 4 months ago

SCPs alone are not sufficient to granting permissions to the accounts in your organization. No permissions are granted by an SCP. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

upvoted 12 times

zhangyu20000 Highly Voted 2 years, 5 months ago

C is correct

SCP policy allow everything except cloudtrail. SCP is boundary but it does not give allow to IAM users. You have to configure allow for every IAM

upvoted 13 times

29fb203 Most Recent 3 months, 3 weeks ago

Selected Answer: A

IAM permissions do not override SCPs. Even if developers have IAM policies allowing s3:CreateBucket, an SCP restriction will still block it unless explicitly allowed.

upvoted 1 times

vmia159 3 months, 3 weeks ago

Your statement is correct but the policy does not deny action on S3. So the SCP is not causing any problems. So it is C.

upvoted 1 times

longlehoang 4 months, 2 weeks ago

Selected Answer: A

{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Action": "*",
"Resource": "*"
},
{
"Sid": "DenyCloudTrail",
"Effect": "Deny",
"Action": "cloudtrail:*",
"Resource": "*"
},
{
"Sid": "AllowS3CreateBucket",
"Effect": "Allow",
"Action": "s3:CreateBucket",
"Resource": "*"
}
]
}

upvoted 1 times

amministrazione 10 months ago

B. Register a block of customer-owned public IP addresses in the AWS account. Create Elastic IP addresses from the address block and assign them to the NAT gateways in the VPC.

upvoted 1 times

helloworldabc 10 months ago

just C

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, SCP is just a distractor, the users need direct permissions

upvoted 3 times

8608f25 1 year, 4 months ago

Selected Answer: C

The problem described does not originate from the Service Control Policy (SCP) itself based on the SCP content provided. The SCP allows all actions ("Action": "") except for actions related to AWS CloudTrail ("Action": "CloudTrail:"), which are explicitly denied. Therefore, the inability for developers to create Amazon S3 buckets is not due to this SCP, as the SCP does not restrict S3 actions.
Given the situation, the correct way to address the developers' inability to create Amazon S3 buckets would be:
* C. Instruct the developers to add Amazon S3 permissions to their IAM entities.

Option C is the correct action because the issue likely stems from the IAM permissions (or lack thereof) assigned to the developers' IAM entities (users, groups, or roles). IAM permissions are required to perform actions within AWS accounts, such as creating S3 buckets. If developers lack the necessary IAM permissions, they would not be able to create S3 buckets regardless of the SCP settings.

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

The SCP in the scenario is allowing any actions with the exception of cloudtrail. Thus, the SCP is not preventing user to create S3 bucket. If the user cannot create a bucket, then the user IAM user/role is missing permissions to create S3 bucket.

upvoted 3 times

shaaam80 1 year, 6 months ago

Selected Answer: C

Answer C.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

it's a C

upvoted 1 times

javitech83 2 years ago

Selected Answer: C

C is correct

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: C

I just wanted to add my vote to the mix to hopefully drown out the wrong votes.
Its definitely C. SCP is only a guardrail, it doesn't actually grant access. So the users would need to be given s3 access separately.
And to address the wrong answer, A isn't correct because creating an s3 bucket is not a cloudtrail action. Being denied cloudtrail wouldn't deny s3 actions.

upvoted 2 times

bhanus 2 years ago

C is the answer. SCP DONT grant permissions. They just set boundaries on what account is capable of giving access to all users. For example, we applied a SCP on an OU that has account A. This SCP has S3fullAWSaccess. This does NOT mean that any IAM user can perform any S3 action. You still need to explicitly define IAM permissions for user to perform action on S3. This is called whitelisting.
Another example, You wrote an SCP that DENIES S3 access and applied it to an OU that has account B. Now Lets say ROOT user of Account B (who got admin previleges) tries to create S3 bucket, they get DENIED error as SCP has already set a bounday saying NOONE in this OU can access S3

upvoted 2 times

Asds 2 years ago

Selected Answer: C

Need to deal with iam policy auth now

upvoted 1 times

Asds 2 years ago

C is right

upvoted 1 times

leehjworking 2 years, 1 month ago

I am not sure the given situation is possible.
When I tested, member (1111-1111-1111) could create bucket without any policy which can be attached or detached by the oneself.

upvoted 2 times

leehjworking 2 years, 1 month ago

Are developers allowed to modify their IAM entities in the situation of option C? If so, I am not sure this is the best practice.

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: C

C is correct

upvoted 2 times

Question #58 *Topic 1*

A company has a monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.

Which solution will meet these requirements?

A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.

B. Create an image of the instance with the reboot option turned on. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

C. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.

D. Create an image of the instance. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Suggested Answer: *A*

Community vote distribution

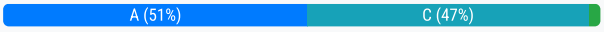

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct answer is C. Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair. Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

Option A is not correct because the instance would need an IAM role with permission to write to S3 and access to the instance via Systems Manager Session Manager.

Option B is not correct because it would require stopping the instance, which would impact the running application.

Option D is not correct because it would require stopping the instance and creating a new EC2 instance, which would impact the running application.

upvoted 38 times

mav3r1ck 1 year, 3 months ago

Not true! Feel free to challenge me if you think I am wrong.

Taking a snapshot of the EBS volume using Amazon DLM is a straightforward approach to ensure data durability and availability. However, this option does not directly address the requirement to move data to an S3 bucket. While EBS snapshots are stored on S3, they are not accessible as regular S3 objects for direct file manipulation or viewing, meaning additional steps would be required to access and use the data in the format specified by the requirement.

Verdict: Does Not Fully Meet Requirements. DLM manages snapshots for EBS volumes but doesn't facilitate direct, accessible backups to S3 as described.

upvoted 11 times

GabrielShiao 8 months, 3 weeks ago

I agree with this A. In addition, the application team has no SSH key access, you can not think that the team has the DLM permission as well. Infrastructure teams generally take this type of role.

upvoted 1 times

ry1999 10 months ago

This is valid, A is the correct answer.

C is wrong because

Explanation: This option indirectly involves copying data to S3. The primary action is taking a snapshot of the EBS volume, which can be

managed by DLM. However, moving the data from a snapshot directly to S3 isn't straightforward. Snapshots are stored in S3 by AWS internally, but this storage is opaque to users and can't be accessed directly as regular S3 objects.

upvoted 2 times

gustori99 1 year, 2 months ago

I'll try to challange you :-)

You can use EBS direct APIs to access data from an EBS snapshot. This is how you can read the data from the snapshot and copy it to S3.

https://docs.aws.amazon.com/ebs/latest/userguide/ebs-accessing-snapshot.html

upvoted 3 times

Sab 1 year, 7 months ago

Your reasoning is wrong . Option A has mentioned that instance profile role is attached to EC2 instance.

upvoted 2 times

Atila50 2 years, 5 months ago

thank you for correcting some of these answers and for the explanations to them

upvoted 3 times

mmendozaf 2 years, 5 months ago

Assuming that EBS is encrypted, I think that is much easier to run the copy command from AW system manager

upvoted 10 times

bititan Highly Voted 2 years, 5 months ago

Selected Answer: A

taking a backup of the data to s3. aws doesn't allow up to view snapshots in s3

upvoted 12 times

tmlong18 1 year, 5 months ago

The requirement is only 'back up'

upvoted 1 times

strike3test Most Recent 1 week ago

Selected Answer: C

This requires that the instance has the SSM agent installed and configured, and the instance profile has the right permissions. However, the question does not confirm that SSM is enabled on the instance.

upvoted 1 times

0dc6cac 2 weeks, 2 days ago

Selected Answer: C

They don't mention whether the instance has SSM permissions, adding that would require a restart of the EC2 instance, causing an interruption. Hence I think C is more appropriate, we can back-up the EBS and there are tools to access data, or we can start up a new instance and get the data when needed.

upvoted 1 times

Kaps443 3 weeks ago

Selected Answer: A

A is the most elegant and operationally safe solution.

It allows you to access the instance securely, without SSH, and back up the data live to S3 — with no service interruption.

upvoted 2 times

eesa 2 months, 1 week ago

Selected Answer: A

Why Option A is correct:

The AWS Systems Manager (SSM) Session Manager allows secure, shell-level access to EC2 instances without the need for an SSH key pair.

Attaching an IAM role with appropriate permissions to the EC2 instance allows it to securely interact with Amazon S3 without manual credential management.

Session Manager does not require a restart of the instance or impact the application availability, ensuring no downtime and continuous service to users.

This solution directly fulfills the requirement: securely copying data from the encrypted EBS volume to S3 without administrative SSH access and without disruption.

upvoted 2 times

grumpysloth 6 months, 2 weeks ago

Selected Answer: D

This is a badly designed question IMO. (D) could be correct but creating an AMI by default will reboot the instance, and no mention of SSM role permissions. (A) could also work but no mention of SSM permissions in the role. Amazon Lnux 2 have pre-installed the SSM agent. (B) is wrong since it interrupts the app. (C) won't work.

upvoted 1 times

Heman31in 6 months, 3 weeks ago

Selected Answer: A

not C because of How EBS Snapshot Export Works
When you export an EBS snapshot to S3, the export creates an Amazon Machine Image (AMI)-compatible format of the snapshot.
This export process results in a snapshot stored as disk image files (e.g., .vmdk, .vhd, .raw, etc.), depending on the format chosen.
The data is not immediately readable or usable as a plain text or object file in S3.
When is the Data Readable?
To make the exported data readable:
Reimport the Snapshot: You would need to reimport the disk image into AWS as a new EBS volume using the VM Import/Export service.
Custom Processing: If the snapshot contains a file system, you could manually process the exported image to extract the data using tools compatible with the format (e.g., mounting a .raw image locally).

upvoted 1 times

Heman31in 6 months, 3 weeks ago

Selected Answer: A

C. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.

Challenges:
While creating an EBS snapshot is feasible without requiring instance access, transferring the data from the snapshot to S3 still requires additional steps.
Snapshot-based backup does not provide direct file-level access for selective backups to S3.
Conclusion: Partially valid, but it does not meet the requirement to back up the data directly to S3. Since it is a legal department ..why to have another copy before transferring to final S3 destination.

upvoted 1 times

Aritra88 6 months, 3 weeks ago

Selected Answer: A

Leverages AWS Systems Manager Session Manager:

Session Manager allows secure shell-less access to the instance without requiring an SSH key.
It provides a way to run commands directly on the instance, even if SSH access is unavailable.
No Disruption to the Application:

The instance remains operational, and the application continues to serve users while the commands are executed.
S3 IAM Role for Access:

By attaching an IAM role to the instance with permissions to write to S3, you can securely transfer data without needing to configure additional credentials.
Efficient and Direct Backup:

Data is copied directly from the running instance to the S3 bucket, eliminating the need for intermediate snapshots, new instances, or additional resources.
Minimal Development Time:

This approach avoids creating images, launching new instances, or performing additional resource management steps.

upvoted 1 times

DhirajBansal 6 months, 4 weeks ago

Selected Answer: A

A is Correct Answer.

IAM Role will provide EC2 instance to write data to S3 bucket.
Systems Manager Session Manager will access system and initiate back writing in S3. This will satisfy the condition of not having SSH Keys.
upvoted 2 times

amministrazione 10 months ago
C. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
upvoted 1 times

Jason666888 11 months ago
Selected Answer: C
Key point: The application must continue to serve the users.
If we choose A, then it may impact the application.
C wouldn't have that problem
upvoted 2 times

Moghite 11 months, 1 week ago
Selected Answer: C
c - Amazon Data Lifecycle Manager allows creation of EBS snapshots
upvoted 3 times

vip2 11 months, 4 weeks ago
Selected Answer: C
C looks more better than A according to keep application running all time
upvoted 2 times

vip2 1 year ago
Selected Answer: C
currect answer is C

Data Lifecycle Manager (DLM) direct APIs can be used to read the data from the snapshot and copy the data to Amazon S3.
upvoted 2 times

cnethers 1 year ago
C
Reason
You can use Amazon Data Lifecycle Manager to automate the creation, retention, and deletion of EBS snapshots and EBS-backed AMIs. When you automate snapshot and AMI management, it helps you to:

Protect valuable data by enforcing a regular backup schedule.

Create standardized AMIs that can be refreshed at regular intervals.

Retain backups as required by auditors or internal compliance.

Reduce storage costs by deleting outdated backups.

Create disaster recovery backup policies that back up data to isolated Regions or accounts.
upvoted 1 times

Question #59 *Topic 1*

A solutions architect needs to copy data from an Amazon S3 bucket m an AWS account to a new S3 bucket in a new AWS account. The solutions architect must implement a solution that uses the AWS CLI.

Which combination of steps will successfully copy the data? (Choose three.)

A. Create a bucket policy to allow the source bucket to list its contents and to put objects and set object ACLs in the destination bucket. Attach the bucket policy to the destination bucket.

B. Create a bucket policy to allow a user in the destination account to list the source bucket's contents and read the source bucket's objects. Attach the bucket policy to the source bucket.

C. Create an IAM policy in the source account. Configure the policy to allow a user in the source account to list contents and get objects in the source bucket, and to list contents, put objects, and set object ACLs in the destination bucket. Attach the policy to the user.

D. Create an IAM policy in the destination account. Configure the policy to allow a user in the destination account to list contents and get objects in the source bucket, and to list contents, put objects, and set objectACLs in the destination bucket. Attach the policy to the user.

E. Run the aws s3 sync command as a user in the source account. Specify the source and destination buckets to copy the data.

F. Run the aws s3 sync command as a user in the destination account. Specify the source and destination buckets to copy the data.

Suggested Answer: *ADF*

Community vote distribution

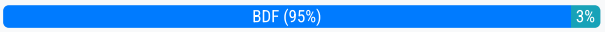

icassp Highly Voted 2 years, 5 months ago

Selected Answer: BDF

"The above command should be executed with destination AWS IAM user account credentials only otherwise the copied objects in destination S3 bucket will still have the source account permissions and won't be accessible by destination account users." According to https://medium.com/tensult/copy-s3-bucket-objects-across-aws-accounts-e46c15c4b9e1.

upvoted 27 times

masetromain 2 years, 5 months ago

You are correct, step E should be executed using the IAM user credentials from the destination account. This is because when objects are copied from one bucket to another, the object's permissions (ACLs) are also copied. Therefore, if the objects are copied using the IAM user credentials from the source account, the objects will have the same permissions as they did in the source bucket, which may not include permissions for the user in the destination account. By using the IAM user credentials from the destination account, the objects will have the appropriate permissions for the user in the destination account once they are copied.

upvoted 5 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: BDF

I switch to BDF;

Step B is necessary so that the user in the destination account has the necessary permissions to access the source bucket and list its contents, read its objects.

Step D is needed so that the user in the destination account has the necessary permissions to access the destination bucket and list contents, put objects, and set object ACLs

Step F is necessary because the aws s3 sync command needs to be run using the IAM user credentials from the destination account, so that the objects will have the appropriate permissions for the user in the destination account once they are copied.

The other choices are not correct because :

A. and C. are about creating policies in the source account but the user who wants to access the data is in the destination account

E. is about running the command with the source account, which is not suitable because it will lead to copied objects in destination S3 bucket still have the source account permissions and won't be accessible by destination account users.

upvoted 16 times

jAtlas7 Most Recent 7 months, 1 week ago

BDF is the answer - see: https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/copy-data-from-an-s3-bucket-to-another-account-and-region-by-using-the-aws-cli.html
upvoted 2 times

amministrazione 10 months ago
B. Create a bucket policy to allow a user in the destination account to list the source bucket's contents and read the source bucket's objects. Attach the bucket policy to the source bucket.
D. Create an IAM policy in the destination account. Configure the policy to allow a user in the destination account to list contents and get objects in the source bucket, and to list contents, put objects, and set objectACLs in the destination bucket. Attach the policy to the user.
F. Run the aws s3 sync command as a user in the destination account. Specify the source and destination buckets to copy the data.
upvoted 1 times

8608f25 1 year, 4 months ago
Selected Answer: BDF
B. Create a bucket policy to allow a user in the destination account to list the source bucket's contents and read the source bucket's objects. Attach the bucket policy to the source bucket. This step ensures that the destination account has the necessary permissions to access the data in the source bucket.
D. Create an IAM policy in the destination account. Configure the policy to allow a user in the destination account to list contents and get objects in the source bucket, and to list contents, put objects, and set object ACLs in the destination bucket. Attach the policy to the user. This step provides the necessary permissions for a user in the destination account to both access the source bucket's contents and write to the destination bucket.
upvoted 1 times

8608f25 1 year, 4 months ago
F. Run the aws s3 sync command as a user in the destination account. Specify the source and destination buckets to copy the data. Performing the sync operation as a user in the destination account, who has been granted the appropriate permissions, ensures that the data can be copied from the source bucket to the destination bucket successfully.
upvoted 1 times

ninomfr64 1 year, 5 months ago
Selected Answer: BDF
Not A. A bucket policy attached to destination bucket cannot allow the source bucket to execute actions
Not C. Because we are picking option B which relies on a policy allowing a user in the destination account.
Not E. Because we are picking options B and D which rely on a user in the destination account
upvoted 1 times

jpa8300 1 year, 5 months ago
Selected Answer: BDF
No need for more explanations, the ones below are enough.
upvoted 1 times

edder 1 year, 7 months ago
Selected Answer: BDF
BD:
https://repost.aws/knowledge-center/cross-account-access-s3
F:
https://docs.aws.amazon.com/cli/latest/userguide/cli-services-s3-commands.html
upvoted 1 times

aviathor 1 year, 10 months ago
Selected Answer: BDF
A is incorrect since a bucket policy cannot allow another bucket to do anything. B. Is however an option since you can indeed create a bucket policy to allow a user in another account to perform operations on the bucket.

Once you have chosen B, then D and F are the only possible choices.
upvoted 2 times

H4des 1 year, 10 months ago
Selected Answer: BCE
BCE should also work
Create bucket policy at destination bucket to allow permission on source aws user
Create IAM policy for source aws user to list/get/put on both buckets
Run s3 sync command from source bucket to destination bucket
upvoted 1 times

CuteRunRun 1 year, 10 months ago

Selected Answer: BDF

I prefer BDF, I do not know why the correct answer is ADF

upvoted 1 times

Christina666 1 year, 11 months ago

Selected Answer: BDF

source bucket: allow destination user + list & get contents permission

destination bucket: allow IAM user to get source bucket contents + destination bucket get/list/put objects + aws sync command

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: BDF

it's BDF for sure

upvoted 1 times

Maria2023 2 years ago

Selected Answer: BDF

The entire idea of A is wrong (you achieve nothing by giving rights from one bucket to another) so we start from B and the rest are a common sense

upvoted 2 times

huanaws088 2 years, 2 months ago

Selected Answer: BDF

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/copy-data-from-an-s3-bucket-to-another-account-and-region-by-using-the-aws-cli.html

upvoted 3 times

God_Is_Love 2 years, 4 months ago

Logical answer : Who ever uploads to a bucket becomes its owner. So A should ring a flaw in it. Similar issue in C. So straight away, A, C are wrong. that points to B,D to be correct. Refer https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/copy-data-from-an-s3-bucket-in-one-account-and-region-to-another-account-and-region.html

Now E or F ? the hint is in D. Destination account user has the necessary privileges to get/put objects permission. So choose destination account or run sync/copy commands. So the answer should be B, D , F

upvoted 6 times

hobokabobo 2 years, 4 months ago

The parts BDF fit together in a way that works.

I think choosing this direction (pulling from the destination account) is slightly more secure than then the other other way round(pushing from source to destination) as only read access is granted to the foreign account but no write access - especially regarding human error: one cannot accidentally tamper with the source, so the worst thing that could happen is that one needs to sync again. The other options don't fit together with other parts.

upvoted 1 times

Question #60 *Topic 1*

A company built an application based on AWS Lambda deployed in an AWS CloudFormation stack. The last production release of the web application introduced an issue that resulted in an outage lasting several minutes. A solutions architect must adjust the deployment process to support a canary release.

Which solution will meet these requirements?

A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.

B. Deploy the application into a new CloudFormation stack. Use an Amazon Route 53 weighted routing policy to distribute the load.

C. Create a version for every new deployed Lambda function. Use the AWS CLI update-function-configuration command with the routing-config parameter to distribute the load.

D. Configure AWS CodeDeploy and use CodeDeployDefault.OneAtATime in the Deployment configuration to distribute the load.

Suggested Answer: *A*

Community vote distribution

A (98%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load is the correct answer as it meets the requirement of supporting a canary release.

Option B is not correct because while it would allow for a canary release, it would involve deploying the new version of the application into a separate CloudFormation stack, which would be a more complex and time-consuming process compared to creating an alias for a new version of the Lambda function.

Option C is not correct because while it would allow for a canary release, it would involve creating a version for every new deployed Lambda function, which would be more complex and time-consuming process compared to creating an alias for a new version of the Lambda function.

upvoted 21 times

masetromain 2 years, 5 months ago

Option D is not correct because AWS CodeDeploy is a deployment service that allows you to automate code deployments to a variety of compute services like EC2 and on-premises servers, but it does not support routing configuration for a canary release on AWS Lambda.

upvoted 6 times

karma4moksha 2 years, 1 month ago

Thank you masetromain, you have been really helpful for taking the time and providing explanation.

upvoted 1 times

Jesuisleon 2 years, 1 month ago

He copied from chatgpt, you didn't find it ?

upvoted 9 times

ninomfr64 1 year, 5 months ago

This is not 100% correct. Actually CodeDeploy support deploy to an AWS Lambda compute platform, the deployment configuration specifies the way traffic is shifted to the new Lambda function versions in your application. You can shift traffic using a canary, linear, or all-at-once deployment configuration. The following lists the predefined configurations available for AWS Lambda canary deployments:

- CodeDeployDefault.LambdaCanary10Percent5Minutes
- CodeDeployDefault.LambdaCanary10Percent10Minutes
- CodeDeployDefault.LambdaCanary10Percent15Minutes
- CodeDeployDefault.LambdaCanary10Percent30Minutes

upvoted 3 times

Jason666888 11 months ago

Yeah the reason D is wrong is not because CodeDeploy doesn't support lambda canary deployment, it's because `OneAtATime` deployment strategy is only for EC2 instances but not for lambdas

upvoted 3 times

Atila50 Highly Voted 2 years, 5 months ago

Selected Answer: A

https://www.examtopics.com/discussions/amazon/view/28312-exam-aws-certified-solutions-architect-professional-topic-1/

upvoted 10 times

29fb203 Most Recent 3 months, 3 weeks ago

Selected Answer: D

AWS CodeDeploy supports canary releases for Lambda functions, which is what the requirement is aiming for.

upvoted 1 times

d401c0d 5 months ago

Selected Answer: A

A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load

upvoted 1 times

Heman31in 6 months, 3 weeks ago

Selected Answer: A

Not D because, the CodeDeployDefault.OneAtATime deployment configuration is primarily designed for EC2 and on-premises instances. For AWS Lambda functions, AWS CodeDeploy provides deployment strategies specific to Lambda, such as Canary, Linear, and All-at-Once.

upvoted 1 times

Aritra88 6 months, 3 weeks ago

Selected Answer: A

Using Lambda Aliases for Canary Releases:

* Lambda aliases are pointers to specific versions of a Lambda function.
* You can use the update-alias command with the routing-config parameter to configure traffic shifting between the current version and the newly deployed version.
* This allows a gradual shift of traffic to the new version while maintaining traffic to the current version.

AWS CLI Example:

* Create a new alias or update an existing alias to shift a portion of the traffic

Testing and Monitoring:

Gradually increase the percentage of traffic to the new version.
Use Amazon CloudWatch metrics to monitor errors, latency, or other performance issues.
Roll back traffic to the previous version if any issues are detected.

```
aws lambda update-alias \
--function-name MyFunction \
--name MyAlias \
--routing-config '{"AdditionalVersionWeights": {"2": 0.10}}'
```

upvoted 2 times

amministrazione 10 months ago

A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.

upvoted 1 times

Chakanetsa 11 months, 1 week ago

Selected Answer: A

A. Create an alias for every new deployed version of the Lambda function. Use the AWS CLI update-alias command with the routing-config parameter to distribute the load.

Explanation:
This option provides a granular level of control for canary deployments:

Versioning: Creating a new version for each deployment ensures that you have a clear record of changes.
Alias: An alias acts as a stable endpoint, allowing you to gradually shift traffic to the new version.

Routing configuration: The routing-config parameter provides fine-grained control over traffic distribution between versions.
By using this approach, you can gradually increase the percentage of traffic to the new version, monitor its performance, and roll back if necessary, minimizing the impact of potential issues.
upvoted 2 times

Chakanetsa 11 months, 1 week ago
Breakdown of other options:
B: While Route 53 weighted routing can distribute traffic, it's less granular than using Lambda aliases and doesn't provide the same level of control.
C: Using the update-function-configuration command doesn't provide the flexibility to gradually shift traffic.
D: CodeDeploy is primarily for deploying code to EC2 instances, not for managing Lambda function traffic.
By using Lambda aliases and the routing-config parameter, you can effectively implement a canary release strategy for your Lambda functions.
upvoted 1 times

ninomfr64 1 year, 5 months ago
Selected Answer: A
Not B. This introduces R53 in the scenario, but we are not sure if R53 fits in the scenario. To combine R53 and Lambda we should use function URL that is not mentioned and we don't know if the app is public. A lot of uncertainty here
Not C. routing-config is an Alias specific configuration aka Weighted Alias and it is not available for the update-function-configuration command https://docs.aws.amazon.com/cli/latest/reference/lambda/update-function-configuration.html
Not D. CodeDeployDefault.OneAtATime is a CodeDeploy option for EC2/on-premise, while in this scenario we need a canary option for Lambda such as CodeDeployDefault.LambdaCanary10Percent5Minutes

A does the job https://docs.aws.amazon.com/cli/latest/reference/lambda/update-alias.html and https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html#configuring-alias-routing
upvoted 1 times

JMAN1 1 year, 6 months ago
100% A is correct. :) I was confused between D and A. But, this url says Codedeployee.AllatOnce deploy option is not for 'canary release'.
https://docs.aws.amazon.com/ko_kr/codedeploy/latest/userguide/deployment-configurations.html
upvoted 2 times

totten 1 year, 9 months ago
Selected Answer: A
Here's why Option A is suitable:

Create an alias: For every new version of your Lambda function, create an alias. Aliases allow you to associate a user-friendly name with a specific version of the function.

Routing configuration: AWS Lambda supports routing configurations that allow you to gradually shift traffic from one alias to another. Using the "routing-config" parameter with the AWS CLI "update-alias" command, you can specify how much traffic each alias should receive.

Gradual release: By configuring the routing, you can control the percentage of traffic directed to the new version (canary). You can gradually increase the traffic percentage as you gain confidence in the new release. If issues arise, you can quickly roll back by adjusting the routing configuration.
upvoted 3 times

Christina666 1 year, 11 months ago
Selected Answer: A
new release-> lambda alias-> update-alias: aws lambda update-alias --function-name my-function --name alias-name --function-version version-number
upvoted 2 times

NikkyDicky 1 year, 12 months ago
Selected Answer: A
D would be an optionn if used Lambda-specific config
upvoted 2 times

SkyZeroZx 2 years ago
Selected Answer: A
keyword = alias for every new deployed version
is a classic usage for deployment canary for lambdas other option usually is codeDeploy but in this options AllAtOnce
then A

upvoted 3 times

AMEJack 2 years, 1 month ago

Sorry OneAtTime

upvoted 1 times

AMEJack 2 years, 1 month ago

Selected Answer: A

CodeDeploy: Although CodeDeploy can help but AllAtOnce is not used for canary traffic shifting.

upvoted 1 times

God_Is_Love 2 years, 4 months ago

Selected Answer: A

aws update-alias command has routing-config option to route the weighted % traffic

As is correct

https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/

Point alias to new version, weighted at 5% (original version at 95% of traffic)

aws lambda update-alias --function-name myfunction --name myalias --routing-config '{"AdditionalVersionWeights" : {"2" : 0.05} }'

upvoted 5 times

Question #61 *Topic 1*

A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.example.com through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

A. Move the EC2 instance into an Auto Scaling group. Place the EC2 instance behind an Application Load Balancer (ALB). Update the DNS record sftp.example.com in Route 53 to point to the ALB.

B. Migrate the SFTP server to AWS Transfer for SFTP. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.

C. Migrate the SFTP server to a file gateway in AWS Storage Gateway. Update the DNS record sftp.example.com in Route 53 to point to the file gateway endpoint.

D. Place the EC2 instance behind a Network Load Balancer (NLB). Update the DNS record sftp.example.com in Route 53 to point to the NLB.

Suggested Answer: *B*

Community vote distribution

tinyflame Highly Voted 2 years, 4 months ago

Selected Answer: B

A=ALB cannot be used with SFTP

B = Correct

C=Storage Gateway is not an SFTP Server

D=NLB can be used with SFTP, but EC2 is single

upvoted 31 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

Option B is the correct answer. Migrating the SFTP server to AWS Transfer for SFTP will improve the reliability and scalability of the SFTP solution. AWS Transfer for SFTP is a fully managed SFTP service that enables the company to transfer files directly into and out of Amazon S3 using the SFTP protocol. By using this service, the company can offload the management of the SFTP server to AWS, which will provide high availability, scalability, and security. The company can then update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname, which will ensure that the SFTP server is reachable on the DNS.

upvoted 14 times

masetromain 2 years, 5 months ago

Option A, C and D do not provide the same level of scalability and reliability as AWS Transfer for SFTP. While placing the EC2 instance behind a load balancer can help improve availability, it will not necessarily improve scalability, and it would still require the company to manage the SFTP server. Option C , migrating the SFTP server to a file gateway in AWS Storage Gateway, would not necessarily improve the scalability and reliability of the SFTP solution, as it would still require the company to manage the SFTP server.

upvoted 4 times

rioisverycute 1 year, 6 months ago

How about the cron job?

upvoted 1 times

amministrazione Most Recent 10 months ago

B. Migrate the SFTP server to AWS Transfer for SFTP. Update the DNS record sftp.example.com in Route 53 to point to the server endpoint hostname.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

B of course

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

keyword = AWS Transfer for SFTP

then B

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: B

B is the way to go..

upvoted 3 times

Question #62 *Topic 1*

A company wants to migrate an application to Amazon EC2 from VMware Infrastructure that runs in an on-premises data center. A solutions architect must preserve the software and configuration settings during the migration.

What should the solutions architect do to meet these requirements?

A. Configure the AWS DataSync agent to start replicating the data store to Amazon FSx for Windows File Server. Use the SMB share to host the VMware data store. Use VM Import/Export to move the VMs to Amazon EC2.

B. Use the VMware vSphere client to export the application as an image in Open Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region. Create and apply an IAM role for VM Import. Use the AWS CLI to run the EC2 import command.

C. Configure AWS Storage Gateway for files service to export a Common Internet File System (CIFS) share. Create a backup copy to the shared folder. Sign in to the AWS Management Console and create an AMI from the backup copy. Launch an EC2 instance that is based on the AMI.

D. Create a managed-instance activation for a hybrid environment in AWS Systems Manager. Download and install Systems Manager Agent on the on-premises VM. Register the VM with Systems Manager to be a managed instance. Use AWS Backup to create a snapshot of the VM and create an AMI. Launch an EC2 instance that is based on the AMI.

Suggested Answer: *D*

Community vote distribution

B (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct answer is B. Use the VMware vSphere client to export the application as an image in Open Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region. Create and apply an IAM role for VM Import. Use the AWS CLI to run the EC2 import command. This approach allows the solutions architect to export the application as an image in OVF format, which preserves the software and configuration settings, and then import it into Amazon EC2 using the EC2 import command.

upvoted 15 times

sammyhaj 1 year, 5 months ago

https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html

upvoted 3 times

masetromain 2 years, 5 months ago

Option A is incorrect because it uses AWS DataSync and FSx for Windows File Server to replicate the data store, but it doesn't preserve the software and configuration settings of the application.

Option C is incorrect because it uses AWS Storage Gateway to export a CIFS share, but it doesn't preserve the software and configuration settings of the application.

Option D is incorrect because it uses AWS Systems Manager and AWS Backup to create a snapshot of the VM, but it doesn't preserve the software and configuration settings of the application.

upvoted 9 times

amministrazione Most Recent 10 months ago

B. Use the VMware vSphere client to export the application as an image in Open Virtualization Format (OVF) format. Create an Amazon S3 bucket to store the image in the destination AWS Region. Create and apply an IAM role for VM Import. Use the AWS CLI to run the EC2 import command.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

A = SMB share cannot host VMware datastore. Also, installing agent modify configuration settings

B = correct

C = not clear how the backup copy is created and what format is used to allow then creating an AMI from it

D = hybrid activation allows SSM to manage on-premise / other cloud VM but doesn't enable AWS Backup. This instead requires a backup gateway to backup VMware environment https://aws.amazon.com/blogs/storage/backup-and-restore-on-premises-vmware-virtual-machines-using-aws-backup/

upvoted 2 times

SorenBendixen 1 year, 10 months ago

Selected Answer: B

The only thing that is missing from the B answer is that the OVF file has to be transformed to a OVA file : https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html.

upvoted 3 times

Brightalw 1 year, 10 months ago

what the B is wrong is that the VM format, should be OVA or VMDK or VHD, not OVF

upvoted 2 times

CuteRunRun 1 year, 10 months ago

Selected Answer: B

I prefer B I do not know why the correct is D.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

it's a B

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: B

https://www.learnitguide.net/2023/01/how-to-migrate-vmware-vm-to-aws-ec2.html

upvoted 3 times

Brightalw 1 year, 10 months ago

It said the VM fomat is OVA or VMDK, not OVF

upvoted 1 times

asifjanjua88 2 years, 2 months ago

I vote to B. Why the admin has selected D as Answer.

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

B is the answer - OVF.

upvoted 2 times

God_Is_Love 2 years, 4 months ago

Selected Answer: B

Use VM Import/Export. B is correct . https://aws.amazon.com/ec2/vm-import/

upvoted 4 times

God_Is_Love 2 years, 4 months ago

https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html

Prerequisites

Create an Amazon S3 bucket for storing the exported images or choose an existing bucket. The bucket must be in the Region where you want to import your VMs. For more information about S3 buckets, see the Amazon Simple Storage Service User Guide.

Create an IAM role named vmimport. For more information, see Required service role.

If you have not already installed the AWS CLI on the computer you'll use to run the import commands, see the AWS Command Line Interface User Guide.

upvoted 2 times

Signup_Nickname 2 years, 5 months ago

Selected Answer: B

I vote B

https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html

upvoted 1 times

Question #63 *Topic 1*

A video processing company has an application that downloads images from an Amazon S3 bucket, processes the images, stores a transformed image in a second S3 bucket, and updates metadata about the image in an Amazon DynamoDB table. The application is written in Node.js and runs by using an AWS Lambda function. The Lambda function is invoked when a new image is uploaded to Amazon S3.

The application ran without incident for a while. However, the size of the images has grown significantly. The Lambda function is now failing frequently with timeout errors. The function timeout is set to its maximum value. A solutions architect needs to refactor the application's architecture to prevent invocation failures. The company does not want to manage the underlying infrastructure.

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

A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).

B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

C. Create an AWS Step Functions state machine with a Parallel state to invoke the Lambda function. Increase the provisioned concurrency of the Lambda function.

D. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of Amazon EC2. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

E. Modify the application to store images on Amazon Elastic File System (Amazon EFS) and to store metadata on an Amazon RDS DB instance. Adjust the Lambda function to mount the EFS file share.

Suggested Answer: *DE*

Community vote distribution

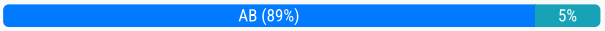

zhangyu20000 Highly Voted 2 years, 5 months ago

A: create Docker image and save it to ECR
B: run this image on Fargate

No answer should have Lambda the will be time out
upvoted 27 times

GabrielShiao 5 months, 1 week ago

While voting A and B, B doesn't really work because Lambda has reached the 15-minute timeout limitation.
upvoted 1 times

masetromain 2 years, 5 months ago

You are correct, both options A and B involve creating a Docker image of the application code and running it on Amazon Elastic Container Service (ECS) using either Fargate or EC2 as the launch type. These options would allow for more control over the resources allocated to the application and potentially prevent timeout errors. Option A is necessary to create the image and store it in a registry, and option B is necessary to run the image on Fargate which is a managed container orchestration service that eliminates the need for provisioning and scaling of the underlying infrastructure.
upvoted 9 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: AB

The correct answer is A and B.

A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).

- This step is necessary to package the application code in a container and make it available for running on ECS.

B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

- This step is necessary to run the containerized application on Fargate, which is a fully managed container orchestration service that eliminates the need to provision and scale the underlying infrastructure.

upvoted 14 times

masetromain 2 years, 5 months ago

Option C and E are not correct because they don't address the problem of timeout errors. AWS Step Functions and Amazon Elastic File System (EFS) are services that can be used to coordinate and manage workflows and file storage respectively, but they don't help with the specific problem of the Lambda function timing out.

Option D is not correct because AWS Fargate is a serverless compute engine for containers that eliminates the need for provisioning and scaling the underlying infrastructure.
It means that the company does not have to manage the underlying infrastructure, which is what the company wants.

upvoted 6 times

amministrazione Most Recent 10 months ago

A. Modify the application deployment by building a Docker image that contains the application code. Publish the image to Amazon Elastic Container Registry (Amazon ECR).
B. Create a new Amazon Elastic Container Service (Amazon ECS) task definition with a compatibility type of AWS Fargate. Configure the task definition to use the new image in Amazon Elastic Container Registry (Amazon ECR). Adjust the Lambda function to invoke an ECS task by using the ECS task definition when a new file arrives in Amazon S3.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: AB

To be honest, I don't trust the examtopic answers anymore, we should only rely on most voted ones

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: AB

AB, ECR + ECS Margate

upvoted 1 times

Ak47g 1 year, 6 months ago

Selected Answer: AB

A: create Docker image and save it to ECR
B: run this image on Fargate

upvoted 1 times

Nicoben 1 year, 6 months ago

Selected Answer: AB

A: create docker image and store in on ECR
B: run it on a AWS-managed infrastructure (as required)

upvoted 1 times

blackgamer 1 year, 8 months ago

The correct answer is A and B. But Lambda function should be replaced with EventBridge.

upvoted 1 times

ggrodskiy 1 year, 8 months ago

Selected Answer: BC

B - 100%
C OR E ??

upvoted 1 times

CuteRunRun 1 year, 10 months ago

Selected Answer: AB

I think is AB

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: AB

it's AB

upvoted 1 times

Jonalb 2 years ago

Selected Answer: AB

AB

its correct!

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: AB

A + B

A , basic dockerized the aplication and use Elastic Container Register

B , deploy how serverless with fargate without overhead managament infrastructure

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

A + B.

upvoted 2 times

dev112233xx 2 years, 3 months ago

Selected Answer: AB

A+B makes sense to me

upvoted 2 times

God_Is_Love 2 years, 4 months ago

Selected Answer: AB

Based on Serverless solutions used, need to go with Fargate in combination with either ECS/EC2.As company does not want to manage infra, we go for because Fargate-ECS combo as Fargate-EC2 needs more maintenance .That means D is out. E is obviously out EFS does not contribute to lambda invocation timeouts.

C is wrong because, increased concurrency (more lambda versions) won't solve timeouts.

That leaves A and B as right answers.

upvoted 4 times

klog 2 years, 4 months ago

Selected Answer: AB

C is not right, question clearly said no involve infrastructure, EC2 is a infrastructure, Lamda time out 15 mins.

upvoted 2 times

Question #64 *Topic 1*

A company has an organization in AWS Organizations. The company is using AWS Control Tower to deploy a landing zone for the organization. The company wants to implement governance and policy enforcement. The company must implement a policy that will detect Amazon RDS DB instances that are not encrypted at rest in the company's production OU.

Which solution will meet this requirement?

A. Turn on mandatory guardrails in AWS Control Tower. Apply the mandatory guardrails to the production OU.

B. Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower. Apply the guardrail to the production OU.

C. Use AWS Config to create a new mandatory guardrail. Apply the rule to all accounts in the production OU.

D. Create a custom SCP in AWS Control Tower. Apply the SCP to the production OU.

Suggested Answer: *B*

Community vote distribution

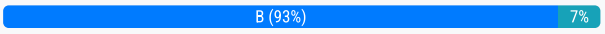

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct answer is B. AWS Control Tower provides a set of "strongly recommended guardrails" that can be enabled to implement governance and policy enforcement. One of these guardrails is "Encrypt Amazon RDS instances" which will detect RDS DB instances that are not encrypted at rest. By enabling this guardrail and applying it to the production OU, the company will be able to enforce encryption for RDS instances in the production environment.

Option A is incorrect because mandatory guardrails are pre-defined by AWS and cannot be customized.
Option C is incorrect because AWS Config does not provide mandatory guardrails for RDS instances.
Option D is incorrect because AWS Control Tower does not provide a feature called custom SCP (Service Control Policy), it uses guardrails instead.

upvoted 20 times

pitakk Highly Voted 2 years, 5 months ago

Selected Answer: B

https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#disallow-rds-storage-unencrypted

upvoted 5 times

Musk 2 years, 5 months ago

The only thing is that this option talks about guardrails, while the article talks about controls, not mandatory.

upvoted 1 times

pk0619 Most Recent 6 months, 1 week ago

Selected Answer: B

Guardrails are now called Controls in Control Tower.

upvoted 1 times

amministrazione 10 months ago

B. Enable the appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower. Apply the guardrail to the production OU.

upvoted 1 times

AloraCloud 12 months ago

The keyword in the question is detect which indicates Config.

"The company must implement a policy that will detect Amazon RDS DB instances that are not encrypted at rest in the company's production OU."

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: B

Option B is correct because AWS Control Tower's strongly recommended guardrails include checks for best practices and additional security measures that are not enforced by default but are highly recommended. Among these, there is likely a guardrail that can detect unencrypted RDS DB

instances, aligning with the company's requirement. Applying this guardrail to the production OU will ensure that all RDS DB instances in that OU are checked for encryption at rest.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

A = Mandatory controls are owned by AWS Control Tower, and they apply by default to every OU on your landing zone and they can't be deactivated
B = correct https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#disallow-rds-storage-unencrypted
C = You cannot create new mandatory controls as they are owned by AWS Control Tower
D = You can create custom SCP in AWS Control Tower as part of the Customizations for AWS Control Tower
https://docs.aws.amazon.com/controltower/latest/userguide/cfcn-set-up-custom-scps.html However this requires a lot of work

upvoted 3 times

ninomfr64 1 year, 5 months ago

Note on D, the question is asking to detect and not to mandate, thus D would not meet requirement

upvoted 3 times

severlight 1 year, 7 months ago

Selected Answer: B

check masetromain's comment

upvoted 1 times

dkx 1 year, 11 months ago

A. No, because mandatory controls are owned by AWS Control Tower, and they apply to every OU on your landing zone. These controls are applied by default when you set up your landing zone, and they can't be deactivated. Moreover, none of them address RDS encrypted at rest.

B. Yes, because Strongly recommended controls are owned by AWS Control Tower. They are based on best practices for well-architected multi-account environments. These controls are not enabled by default, and they can be deactivated through the AWS Control Tower console or the control APIs. Moreover, three of them are RDS detective controls

C. No, because AWS Config does not create mandatory guardrails; AWS Config has managed and custom rules

D. No, because SCPs are created in AWS Orgs and are not designed to detect Amazon RDS DB instances that are not encrypted at rest.

upvoted 4 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

It's. B

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

A seems but previous exist rule
then B is more apropiate in this case
https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#disallow-rds-storage-unencrypted

upvoted 1 times

EricZhang 2 years, 1 month ago

C - using AWS Config for detective action

upvoted 2 times

OCHT 2 years, 2 months ago

Selected Answer: C

Option B suggests enabling an appropriate guardrail from the list of strongly recommended guardrails in AWS Control Tower and applying it to the production OU. While AWS Control Tower provides a set of pre-packaged guardrails that enforce best practices for security, operations, and compliance, there is no guarantee that there is a pre-packaged guardrail specifically for detecting Amazon RDS DB instances that are not encrypted at rest.

In contrast, option C creates a custom rule in AWS Config that specifically checks for Amazon RDS DB instances that are not encrypted at rest. This provides more flexibility and control in ensuring that the company's specific requirement is met.

upvoted 3 times

passthatexam1 2 years, 2 months ago

It's incorrect ideally you only apply to the OU and not to an individual account, therefore this needs to be discounted.

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

Enable the appropriate guardrail

upvoted 2 times

Ajani 2 years, 3 months ago

Selected Answer: B

Mandatory controls are owned by AWS Control Tower, and they apply to every OU on your landing zone. These controls are applied by default when you set up your landing zone, and they can't be deactivated.

The solution requirement falls under a proactive(Recommended Control).

https://docs.aws.amazon.com/controltower/latest/userguide/rds-rules.html#ct-rds-pr-16-description

Optional controls are OU specific.

upvoted 4 times

God_Is_Love 2 years, 4 months ago

Selected Answer: B

Tip - As this detective guardrail is available, answer is B. But if the guardrail is not available in that predefined list, the answer would be --C

https://aws.amazon.com/blogs/mt/aws-control-tower-detective-guardrails-as-an-aws-config-conformance-pack/

upvoted 3 times

klog 2 years, 4 months ago

Selected Answer: B

question is asking for detection, not mandate

upvoted 2 times

Question #65 *Topic 1*

A startup company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI. The company's engineers rely heavily on SSH access to the instances for troubleshooting.

The company's existing architecture includes the following:

• A VPC with private and public subnets, and a NAT gateway.
• Site-to-Site VPN for connectivity with the on-premises environment.
• EC2 security groups with direct SSH access from the on-premises environment.

The company needs to increase security controls around SSH access and provide auditing of commands run by the engineers.

Which strategy should a solutions architect use?

A. Install and configure EC2 Instance Connect on the fleet of EC2 instances. Remove all security group rules attached to EC2 instances that allow inbound TCP on port 22. Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI.

B. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices. Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs.

C. Update the EC2 security groups to only allow inbound TCP on port 22 to the IP addresses of the engineer's devices. Enable AWS Config for EC2 security group resource changes. Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules.

D. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attached. Attach the IAM role to all the EC2 instances. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.

Suggested Answer: *D*

Community vote distribution

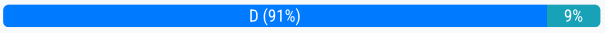

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: D

The correct answer is D. This strategy uses IAM roles and AWS Systems Manager to provide secure and auditable SSH access to the instances. The IAM role is attached to all the EC2 instances and has the AmazonSSMManagedInstanceCore managed policy attached, which allows the instances to be managed by Systems Manager. The engineers then install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager. This approach provides secure and auditable access to the instances without the need for IP-based security group rules or additional infrastructure.

upvoted 20 times

masetromain 2 years, 5 months ago

Option A uses EC2 Instance Connect to provide secure and auditable SSH access to the instances, but it requires additional infrastructure and configuration.

Option B provides auditing of commands run by the engineers, but it relies on IP-based security group rules, which can be difficult to manage and may not be as secure as using IAM roles.

Option C uses AWS Config and Firewall Manager to automatically remediate changes to security group rules, but it still relies on IP-based security group rules and does not provide an auditable method of access to the instances.

upvoted 4 times

masetromain 2 years, 5 months ago

For option A to work, the following additional infrastructure and configuration would be required:

The EC2 Instance Connect service needs to be enabled in the AWS account and the appropriate IAM permissions would need to be granted to the engineers.

The EC2 instances would need to have the EC2 Instance Connect agent installed and configured.

The engineers would need to install the EC2 Instance Connect CLI on their devices and have the necessary credentials to authenticate with AWS.

In addition, the company would need to update their processes and procedures to ensure that engineers are only using EC2 Instance Connect to access the instances and that all access is being logged and audited.
upvoted 4 times

adrian202 1 year, 6 months ago
The key factor is that Option A explains to remove the port 22 inbound SSH access security group, they would need to keep that present: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-connect-prerequisites.html
upvoted 3 times

God_Is_Love Highly Voted 2 years, 3 months ago
Selected Answer: D
A is wrong because Instance connect does not provided auditing
B is wrong because it mentions OS audit logs. we need to audit SSH trafic
C is wrong because we want to audit not remediate as asked in question. config service is to record
using predefined rules and remediate as well

D is correct because,
By attaching the AmazonSSMManagedInstanceCore policy to an IAM role, EC2 instances can be controlled and monitored through the Systems Manager service, enabling capabilities such as remote instance management, patching, and compliance reporting. (ChatGPT response its answers are brief and helpful sometimes)
upvoted 11 times

kgpoj 9 months, 3 weeks ago
The explanation for A is wrong.
AWS EC2 Instance Connect does support auditing.
upvoted 1 times

amministrazione Most Recent 10 months ago
D. Create an IAM role with the AmazonSSMManagedInstanceCore managed policy attached. Attach the IAM role to all the EC2 instances. Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin for their devices and remotely access the instances by using the start-session API call from Systems Manager.
upvoted 1 times

gofavad926 1 year, 3 months ago
Selected Answer: D
D, use SSM
upvoted 1 times

8608f25 1 year, 4 months ago
Selected Answer: D
Option D is the best strategy because it leverages AWS Systems Manager Session Manager, which allows for secure instance management without the need for SSH access. By attaching an IAM role with the AmazonSSMManagedInstanceCore policy to EC2 instances, engineers can use Session Manager for shell access to instances without needing to open port 22, significantly enhancing security. Session Manager also automatically logs session activity to S3 or CloudWatch Logs, providing the required command auditing capability. This eliminates the need for direct SSH access and offers a centralized, secure, and audited method for engineers to access and run commands on instances.
upvoted 2 times

rioisverycute 1 year, 6 months ago
Selected Answer: B
It required to increase security around ssh access, why so many people voted on D?
upvoted 1 times

djeong95 1 year, 4 months ago
Cloudwatch agent does not provide auditable logs for SSH sessions; it only provides metrics about CPU/Memory/Network Packets/etc; nothing about what user started session at what time and ran certain trackable API calls while in that session.
upvoted 1 times

Chung234 1 year, 8 months ago

The answer is D. Option A is wrong because EC2 Instance Connect requires the host security group to permit SSH traffic inbound.
https://repost.aws/questions/QUnV4R9EoeSdW0GT3cKBUR7w/what-is-the-difference-between-ec2-instance-connect-and-session-manager-ssh-connections

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

It's D

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: D

keyword = AWS Systems Manager Session Manager
then D

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: D

D for sure.

upvoted 2 times

Ajani 2 years, 3 months ago

Why its NOT A
To connect using the Amazon EC2 console, the instance must have a public IPv4 address.

If the instance does not have a public IP address, you can connect to the instance over a private network using an SSH client or the EC2 Instance Connect CLI. For example, you can connect from within the same VPC or through a VPN connection, transit gateway, or AWS Direct Connect.

EC2 Instance Connect does not support connecting using an IPv6 address.
going with D:

upvoted 2 times

lygf 2 years, 4 months ago

Selected Answer: D

Need to be able to audit the commands ran on the machine.

upvoted 2 times

DWsk 2 years, 4 months ago

I don't understand why it can't be A for this one. Why is AWS Systems Manager Session better than EC2 Instance Connect? They both require installing something on the instances.

upvoted 1 times

lygf 2 years, 4 months ago

Could option A audit the commands ran on the server, as required by the question? I knew D certainly can.

upvoted 1 times

anita_student 2 years, 3 months ago

For EC2 instance connect there are a few requirements:
- instance has public IP (the instances in question are private)
- you have port 22 open (A says remove port 22 inbound)

upvoted 4 times

moota 2 years, 4 months ago

Selected Answer: D

According to ChatGPT,

Yes, AWS Systems Manager Session Manager can track the commands that are executed during a session. The session is recorded in the form of a log, which can be accessed and reviewed later. The log contains information such as the start time, end time, and the user who initiated the session, as well as a record of all the commands executed during the session, including their output and exit codes. This information can be useful for auditing purposes, troubleshooting, and compliance reporting.

upvoted 2 times

tinyflame 2 years, 4 months ago

Selected Answer: B

provide auditing of commands run by the engineers = B Only

upvoted 3 times

joefromnc 1 year, 10 months ago

Read docs you can audit command using SSM https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager-logging.html

upvoted 1 times

rlf 1 year, 8 months ago

"In addition to providing information about current and completed sessions in the Systems Manager console, Session Manager provides you with the ability to audit session activity in your AWS account using AWS CloudTrail"

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

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

upvoted 1 times

Question #66 *Topic 1*

A company that uses AWS Organizations allows developers to experiment on AWS. As part of the landing zone that the company has deployed, developers use their company email address to request an account. The company wants to ensure that developers are not launching costly services or running services unnecessarily. The company must give developers a fixed monthly budget to limit their AWS costs.

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

A. Create an SCP to set a fixed monthly account usage limit. Apply the SCP to the developer accounts.

B. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process.

C. Create an SCP to deny access to costly services and components. Apply the SCP to the developer accounts.

D. Create an IAM policy to deny access to costly services and components. Apply the IAM policy to the developer accounts.

E. Create an AWS Budgets alert action to terminate services when the budgeted amount is reached. Configure the action to terminate all services.

F. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reached. Invoke an AWS Lambda function to terminate all services.

Suggested Answer: *BDF*

Community vote distribution

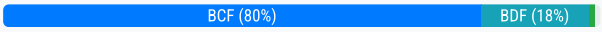

spd Highly Voted 2 years, 4 months ago

Selected Answer: BCF

Clear - BCF - SCP is preferable over IAM

upvoted 20 times

kiran15789 Highly Voted 2 years, 4 months ago

Selected Answer: BCF

I prefer D over C as IAM cant be applied to Account

upvoted 16 times

AWSum1 9 months, 2 weeks ago

Option D says apply it to the Developers accounts. Unnecessary operational overhead

upvoted 1 times

vda2024 Most Recent 7 months, 3 weeks ago

Fixed monthly budget> implement AWS budget. hence B is correct.

Prevent running unnecessary services> implement SCP. hence C is correct.

on F: I'm just not sure why do we want to terminate all resources and why not just don't let them run additional.

upvoted 2 times

amministrazione 10 months ago

B. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process.

C. Create an SCP to deny access to costly services and components. Apply the SCP to the developer accounts.

F. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reached. Invoke an AWS Lambda function to terminate all services.

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

Selected Answer: BCF

Why Option C is Preferred to D :

Centralized Control: SCPs provide a centralized way to manage permissions across all accounts in an organization, ensuring consistent enforcement of policies.

Scalability: SCPs are easier to manage and scale when dealing with multiple accounts, as changes to the SCP will automatically apply to all accounts under the organization.

Compliance: SCPs help ensure compliance with organizational policies by preventing the use of restricted services across all accounts.

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: BCF

BCF - SCP, budget and custom lambda to terminate services

upvoted 2 times

wooin992 1 year, 3 months ago

Selected Answer: BDF

BDF

cannot apply scp in account, need to apply it in OU

upvoted 1 times

MAZIADI 10 months, 3 weeks ago

wrong, you can apply scp to an account

upvoted 2 times

8608f25 1 year, 4 months ago

Selected Answer: BCF

B. Use AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process. AWS Budgets allows you to set custom cost and usage budgets that alert you when you exceed your thresholds.

C. Create an SCP to deny access to costly services and components. Apply the SCP to the developer accounts. By creating an SCP that specifically denies access to costly AWS services, the company can prevent developers from launching such services, thereby helping to keep costs within the fixed monthly budget.

F. Create an AWS Budgets alert action to send an Amazon Simple Notification Service (Amazon SNS) notification when the budgeted amount is reached. Invoke an AWS Lambda function to terminate all services. While AWS Budgets cannot directly terminate services when a budget is exceeded, you can configure an alert to trigger a notification. This notification can then invoke a Lambda function designed to assess and terminate services as necessary, based on the company's policies.

upvoted 2 times

duriselvan 1 year, 5 months ago

Setting a monthly cost budget with a variable target amount, with each subsequent month growing the budget target by 5 percent. Then, you can configure your notifications for 80 percent of your budgeted amount and apply an action. For example, you could automatically apply a custom IAM policy that denies you the ability to provision additional resources within an account.

https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-managing-costs.html

ans :bdf

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: BCF

A = SCP is used to limit permission that administrator can grant IAM users/roles, SCP cannot set a fixed monthly account usage limit

B = correct

C = correct

D = it could work, but it would required more work wrt SCP

E = Budget actions cannot terminate all kind of services, actually supports 3 types of actions 1/ apply IAM policy to IAM identities, 2/ apply SCP to an OU and 3/ terminate EC2 and RDS instances

F = correct

upvoted 3 times

jpa8300 1 year, 5 months ago

Selected Answer: BDF

Although, C is correct, some people here says that SCP cannot be attached to an account, but it is not true, you can, the most common option when we want to deny permissions to an account is to use an IAM policy.

upvoted 1 times

rlf 1 year, 8 months ago

BCF.

In Option D, we can not apply IAM policy to an AWS Account.

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: BDF

I'd go with BDF, since there's no mention of OU. As a rule of thumb, IAM policies to restrict are applied on Accounts, Users, Groups and SCP's on OU's.

upvoted 4 times

vn_thanhtung 1 year, 10 months ago

IAM policies for user ? https://docs.aws.amazon.com/kms/latest/developerguide/iam-policies-overview.html

upvoted 1 times

vn_thanhtung 1 year, 10 months ago

Sorry I mistake, IAM policies can applied on User.

upvoted 1 times

CuteRunRun 1 year, 10 months ago

Selected Answer: BCF

BCF is right.

I think SCP is more convenient than iam.

You need to config the IAM to all account manually

upvoted 2 times

[Removed] 1 year, 11 months ago

Selected Answer: BCF

prefer SCP over IAm in org accounts

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: BCF

It's a BCF

upvoted 2 times

PhuocT 2 years ago

Selected Answer: BCF

C - SCP would be prefer to control the services could be used in Organization's AWS accounts.

upvoted 2 times

Question #67 *Topic 1*

A company has applications in an AWS account that is named Source. The account is in an organization in AWS Organizations. One of the applications uses AWS Lambda functions and stores inventory data in an Amazon Aurora database. The application deploys the Lambda functions by using a deployment package. The company has configured automated backups for Aurora.

The company wants to migrate the Lambda functions and the Aurora database to a new AWS account that is named Target. The application processes critical data, so the company must minimize downtime.

Which solution will meet these requirements?

A. Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

B. Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager {AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.

C. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions and the Aurora DB cluster with the Target account. Grant the Target account permission to clone the Aurora DB cluster.

D. Use AWS Resource Access Manager (AWS RAM) to share the Lambda functions with the Target account. Share the automated Aurora DB cluster snapshot with the Target account.

Suggested Answer: *C*

Community vote distribution

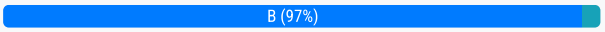

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct answer is option B. This solution uses a combination of AWS Resource Access Manager (RAM) and automated backups to migrate the Lambda functions and the Aurora database to the Target account while minimizing downtime.

In this solution, the Lambda function deployment package is downloaded from the Source account and used to create new Lambda functions in the Target account. The Aurora DB cluster is shared with the Target account using AWS RAM and the Target account is granted permission to clone the Aurora DB cluster, allowing for a new copy of the Aurora database to be created in the Target account. This approach allows for the data to be migrated to the Target account while minimizing downtime, as the Target account can use the cloned Aurora database while the original Aurora database continues to be used in the Source account.

upvoted 21 times

masetromain 2 years, 5 months ago

Option A is not the best solution because it doesn't share the Aurora DB cluster with the Target account and this would cause data inconsistencies as the Source and Target accounts would not share the same data.

Option C is not the best solution because, it does not specify how the data will be migrated and it would cause downtime as the Source and Target accounts are not sharing the same data.

Option D is not the best solution because it does not specify how the Lambda function will be migrated and it would cause data inconsistencies as the Source and Target accounts are not sharing the same data.

upvoted 2 times

lxrdm 1 year, 12 months ago

For option A, its also not possible because automated snapshots cannot be shared..
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-share-snapshot.html

upvoted 4 times

Simon523 Highly Voted 1 year, 10 months ago

Selected Answer: B

AWS Resource Access Manager (RAM) can only share the follow services:
- Amazon Aurora – DB clusters

- Amazon EC2 – capacity reservations and dedicated hosts
- AWS License Manager – License configurations
- AWS Outposts – Local gateway route tables, outposts, and sites
- Amazon Route 53 – Forwarding rules
- Amazon VPC – Customer-owned IPv4 addresses, prefix lists, subnets, traffic mirror targets, transit gateways, transit gateway multicast domains

https://docs.aws.amazon.com/ram/latest/userguide/shareable.html

upvoted 18 times

JOJO9 Most Recent 6 months, 2 weeks ago

Selected Answer: A

Lambda function and Aurora cluster can NOT be shared with RAM!

upvoted 2 times

amministrazione 10 months ago

B. Download the Lambda function deployment package from the Source account. Use the deployment package and create new Lambda functions in the Target account. Share the Aurora DB cluster with the Target account by using AWS Resource Access Manager {AWS RAM). Grant the Target account permission to clone the Aurora DB cluster.

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: B

A is viable, but as AWS RAM can share Aurora clusters, B is faster. However, AWS RAM can't share lambdas, so C and D are out.

upvoted 3 times

mnsait 7 months, 1 week ago

Option A is NOT viable. As @lxrdm has pointed out with documentation, it is not possible to share 'automated' db cluster snapshots.

upvoted 1 times

Dgix 1 year, 3 months ago

B, C, and D are all out since AWS RAM cannot share either Lambdas or Aurora DB clusters. A is the only viable one - you must use a manual shapshot for the DB, share it, and redeploy any deployment package in the destination account. (The question tries to trip you up by its wording: lambda deployments can't be downloaded, but the same deployment packages used to deploy the lambdas can, for instance from S3 or from source)

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: B

Option B is the most accurate and efficient solution based on this AWS article content (https://aws.amazon.com/about-aws/whats-new/2019/07/amazon_aurora_supportscloningacrossawsaccounts-/). It correctly outlines the steps for Lambda migration and utilizes the Aurora DB cluster cloning feature across accounts via AWS RAM, which aligns with the article's description. This approach ensures minimal downtime and efficient migration by allowing direct cloning of the Aurora database.

Option C incorrectly suggests using AWS RAM to share Lambda functions, which is not supported yet based on latest sharable AWS resources: https://docs.aws.amazon.com/ram/latest/userguide/shareable.html

upvoted 2 times

master9 1 year, 5 months ago

Selected Answer: C

AWS Resource Access Manager (RAM) to share AWS Lambda functions and Aurora DB clusters with another AWS account. AWS RAM allows you to share resources that are created and managed by other AWS services with individual AWS accounts or with the accounts in an organization or organizational units (OUs) in AWS Organizations.

To share a Lambda function with another AWS account, you can delegate access to an IAM user (or all users) in the other AWS account so that they can assume a role in your account and invoke the Lambda function in your account.

To share an Aurora DB cluster with another AWS account, you can create a resource share in AWS RAM and specify the Amazon Resource Name (ARN) of the Aurora DB cluster as the resource to share. You can then specify the AWS account IDs of the accounts with which you want to share the resource.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

A = you can share snapshot to restore DB, but this will introduce some downtime

B = correct (cloning a DB allows for very limited downtime)

C = if you only share Lambda you are not migrating it, also it appears the Lambda is not a RAM sharable resource
https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html
D = it appears the Lambda is not a RAM sharable resource and you cannot directly share an automated snapshot, you need first to create a manual snapshot by copying the automated snapshot, and then share that copy https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-share-snapshot.html
upvoted 1 times

ninomfr64 1 year, 5 months ago
A is not correct as you cannot directly share an automated snapshot, you need first to create a manual snapshot by copying the automated snapshot, and then share that copy https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-share-snapshot.html
upvoted 1 times

learnwithaniket 1 year, 6 months ago
Selected Answer: B
There is limit on the number of resources you can share with AWS RAM.
AWS RAM does not support direct sharing of Lambda functions between accounts.
https://docs.aws.amazon.com/ram/latest/userguide/shareable.html
upvoted 2 times

NikkyDicky 1 year, 12 months ago
Selected Answer: B
it's B.
In A - automated snapshots are not shareable
upvoted 2 times

Maria2023 2 years ago
Selected Answer: B
Option B minimizes downtime, compared to A, where we only share a snapshot of the cluster. For C we do not migrate the lambdas, we just share them, which is not the idea of the exercise.
upvoted 1 times

SkyZeroZx 2 years ago
Selected Answer: B
The correct answer is option B. This solution uses a combination of AWS Resource Access Manager (RAM) and automated backups to migrate the Lambda functions and the Aurora database to the Target account while minimizing downtime.
upvoted 1 times

SkyZeroZx 2 years ago
in case the letter A use only snapshot not sync the complete data and is posible lost data in the process
upvoted 2 times

Perkuns 2 years ago
Selected Answer: C
They just want to migrate the Lambda and Aurora DB, they dont care about the app itself
upvoted 1 times

rbm2023 2 years, 1 month ago
Selected Answer: B
The question is about migration and not sharing, so the answer is how to use a RAM feature to help you on the migration. In option D they are not migrating anything, both Lambda and Aurora are being shared with the Target account and not migrated. In option C is a similar situation, the Lambda is not being migrated. Option A seems a good option but might cause a larger downtime. Hence option D is more appropriate because you can use the cluster share with the Target account and clone the database cluster into it. In my view this answer should contemplate in which moment the cutoff from Source to Target would occur.
upvoted 3 times

takecoffe 2 years, 2 months ago
Selected Answer: B
You can share the following Amazon Aurora resources by using AWS RAM.
upvoted 2 times

mfsec 2 years, 3 months ago
Selected Answer: B
B is the way forward
upvoted 2 times

Question #68 *Topic 1*

A company runs a Python script on an Amazon EC2 instance to process data. The script runs every 10 minutes. The script ingests files from an Amazon S3 bucket and processes the files. On average, the script takes approximately 5 minutes to process each file The script will not reprocess a file that the script has already processed.

The company reviewed Amazon CloudWatch metrics and noticed that the EC2 instance is idle for approximately 40% of the time because of the file processing speed. The company wants to make the workload highly available and scalable. The company also wants to reduce long-term management overhead.

Which solution will meet these requirements MOST cost-effectively?

A. Migrate the data processing script to an AWS Lambda function. Use an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects.

B. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure Amazon S3 to send event notifications to the SQS queue. Create an EC2 Auto Scaling group with a minimum size of one instance. Update the data processing script to poll the SQS queue. Process the S3 objects that the SQS message identifies.

C. Migrate the data processing script to a container image. Run the data processing container on an EC2 instance. Configure the container to poll the S3 bucket for new objects and to process the resulting objects.

D. Migrate the data processing script to a container image that runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Create an AWS Lambda function that calls the Fargate RunTaskAPI operation when the container processes the file. Use an S3 event notification to invoke the Lambda function.

Suggested Answer: *D*

Community vote distribution

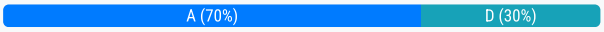

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

The correct answer is A, migrating the data processing script to an AWS Lambda function and using an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects. This solution meets the company's requirements of high availability and scalability, as well as reducing long-term management overhead, and is likely to be the most cost-effective option.

Option B involves creating an SQS queue and configuring S3 to send event notifications to it. The data processing script would then poll the SQS queue and process the S3 objects that the SQS message identifies. While this option also provides high availability and scalability, it is less cost-effective than using Lambda, as it requires additional resources such as an SQS queue and an EC2 Auto Scaling group.

upvoted 22 times

hamimelon 1 year, 9 months ago

Agree. Also, it says the company does not wanna manage long-term overhead, which points to serverless.

upvoted 3 times

dpatra 1 year, 8 months ago

SQS is out of the question because the script already has a built in logic that will prevent it to reprocess a message that's already been processed

upvoted 1 times

masetromain 2 years, 5 months ago

Option C, migrating the data processing script to a container image and running it on an EC2 instance, would still require the company to manage the underlying EC2 instances and may not be as cost-effective as using Lambda.

Option D, migrating the data processing script to a container image that runs on Amazon ECS on AWS Fargate, would still require the company to manage the underlying infrastructure and may not be as cost-effective as using Lambda. Additionally, it introduces additional complexity by adding a Lambda function that calls the Fargate RunTask API operation.

upvoted 5 times

red_panda 1 year, 3 months ago

ECS in Fargate mode you don't need to manage anything underling infra!
You're totally forgot about cost, for sure running an ECS Fargate has lower cost than running a Lambda for 5 minutes every 10 minutes!
Also the function to trigger the ECS workload (in option D), running for milliseconds (as need only to notify the doc upload in S3), so it's more correct the D answer.
Ask to any Gen AI model, you will have mine answer with more details :)
upvoted 1 times

NirvanaSNM 11 months, 2 weeks ago
Use an S3 event notification to invoke the Lambda function to process the objects
upvoted 1 times

zhangyu20000 Highly Voted 2 years, 5 months ago
A is correct, it provide HA, scale, less management. Task only need 5 minutes
B: enen more complex
C: container still run on one EC2, not scale
d: need container, Farget and Lambda. Complex than A
upvoted 7 times

RB100 Most Recent 4 weeks, 1 day ago
Selected Answer: A
The Lambda solution (Option A) provides the most efficient, cost-effective, and manageable solution while meeting all requirements for high availability and scalability with minimal operational overhead.
upvoted 1 times

albert_kuo 3 months, 4 weeks ago
Selected Answer: A
+---------------------+
| Amazon S3 |
| (Stores Uploaded Files)|
+---------------------+
|
| (S3 Event Notification)
v
+---------------------+
| AWS Lambda |
| (Processes Files) |
+---------------------+
|
| (Processed Data)
v
+---------------------+
| Amazon S3 |
| (Stores Processed Data)|
+---------------------+
upvoted 1 times

SIJUTHOMASP 7 months, 1 week ago
Option with Lambda would be more reasonable because the rational behind the cost on the solution would be triggering lambda on the S3 event. The current behaviour is 40% of EC2 being not utilised so that Option D to run in ECS with Fargate would be costlier than Lambda option here
upvoted 1 times

amministrazione 10 months ago
A. Migrate the data processing script to an AWS Lambda function. Use an S3 event notification to invoke the Lambda function to process the objects when the company uploads the objects.
upvoted 1 times

MAZIADI 10 months, 3 weeks ago
Selected Answer: A
instead of scheduled 10 min with multiple files processing (each takes 5 minutes) it will be event driven with lambda each time a file is uploaded --> Answer A
upvoted 1 times

zolthar_z 11 months, 2 weeks ago

Selected Answer: A

A, the company wants to reduce management overhead not costs, we should stay with the question requirement, it doesn't said anything about cost, probably D will be cheaper but the solution must resolve the question necessity and is reduce long-term management overhead

upvoted 2 times

Helpnosense 1 year ago

Selected Answer: D

I vote D because the service is from EC2 to lambda and work is processing data. Without given how big is the data we can't assume that the data is always below the lambda ephemeral storage limit 0.5GB. Nowadays, a file can easily break 0.5GB.
While D is still EC2 based so whatever previous EC2 can do farget can do as well.

upvoted 1 times

Shenannigan 1 year ago

Selected Answer: A

The answer is A:

AWS Pricing Calculator
(using:
10,000 request per month,
300,000 ms which = 5 minutes
128 MB of Memory
512 MB of Storage
)

Amount of memory allocated: 128 MB x 0.0009765625 GB in a MB = 0.125 GB
Amount of ephemeral storage allocated: 512 MB x 0.0009765625 GB in a MB = 0.5 GB
Pricing calculations
10,000 requests x 300,000 ms x 0.001 ms to sec conversion factor = 3,000,000.00 total compute (seconds)
0.125 GB x 3,000,000.00 seconds = 375,000.00 total compute (GB-s)
375,000.00 GB-s x 0.0000166667 USD = 6.25 USD (monthly compute charges)
10,000 requests x 0.0000002 USD = 0.00 USD (monthly request charges)
0.50 GB - 0.5 GB (no additional charge) = 0.00 GB billable ephemeral storage per function
Lambda costs - Without Free Tier (monthly): 6.25 USD

For those thinking D is the cheaper option, do you really believe ECS Fargate would be cheaper?

upvoted 2 times

red_panda 1 year, 3 months ago

Selected Answer: D

Ok i was thinking between A and D.
I'm pretty sure which is D our answer, see the details.

The requirements are:
- COST as much as possible low
- OPERATIONS as much as possible managed.

So at the first reading, the A option seems to be the correct option (because it's totally AWS managed), but here we're totally forgot the cost.
Running a Lambda function, for 5 minutes every 10 minutes, it's very very more expensive than a simple ECS task running continously.

Finally, ECS in fargate mode is totally AWS managed, so we will have lower cost, and a serverless and HA environment, which auto-scale if we need more processing at time.

For me, option D is the correct answer.

upvoted 3 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A, use lambda function is much cost-effective than use ECS Margate

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: A

Option A is the most cost-effective and efficient solution. AWS Lambda allows for running code in response to triggers such as S3 event notifications without the need to manage servers, thereby directly addressing the requirement to reduce long-term management overhead. Since the script is only needed when new files are uploaded and takes about 5 minutes to process each file, Lambda's ability to scale automatically and its billing model based on actual compute time used make it an ideal solution. Lambda can process files immediately upon upload, maximizing efficiency and minimizing idle time.

Option D proposes using Amazon ECS on AWS Fargate with Lambda to trigger tasks. This solution introduces container orchestration, which can improve scalability and reduce some management overhead. However, it is not as cost-effective as directly invoking a Lambda function to process files, considering the lightweight nature of the task and the added complexity of managing container orchestration and Lambda functions together.

upvoted 1 times

LazyAutonomy 1 year, 4 months ago

Selected Answer: D

100% the answer is D.

5 minutes to process EACH FILE? And the EC2 instance is processing files 60% of the time?

Lambda would be crazy expensive in this scenario. ECS/Fargate = cheaper for sure. See link in @covabix879 comment for proof of this.

Greyeye said something rather ridiculous: "If you get 1000 images, you will see 1000 tasks. That is not economical or cheap."

How can 1x EC2 instance running a script every 10 minutes process 1000 images with each one taking 5 minutes? Even if the script processed images in parallel, e.g. one image per vCPU at a time, that instance would need 500 vCPUs! For the EC2 instance to be idle 40% of the time, it would need 833 vCPUs. That's ridiculous.

But even if 1000 images suddenly appeared, the Lambda solution would still result in 1000 Lambdas all firing and running for 5 minutes each. Which is going to be more expensive than ECS/Fargate.

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: A

A = correct

B = does not reduce long-term management overhead

C = does not reduce long-term management overhead

D = does not reduce long-term management overhead

Note: D is a cheap options as mentioned by other here below could be cheaper than A. However, in addition to maintaining the script code it requires to maintain the container image and the lambda

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: A

in the real world it might be D, but with provided details and keeping in mind lambda retries in case of A, I would vote for A.

upvoted 1 times

Sandeep_B 1 year, 8 months ago

Selected Answer: A

D is more complex and overload for administration. Hence Vote for A

upvoted 1 times

Question #69 *Topic 1*

A financial services company in North America plans to release a new online web application to its customers on AWS. The company will launch the application in the us-east-1 Region on Amazon EC2 instances. The application must be highly available and must dynamically scale to meet user traffic. The company also wants to implement a disaster recovery environment for the application in the us-west-1 Region by using active-passive failover.

Which solution will meet these requirements?

A. Create a VPC in us-east-1 and a VPC in us-west-1. Configure VPC peering. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs. Place the Auto Scaling group behind the ALB.

B. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPC. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPC. Place the Auto Scaling group behind the ALSet up the same configuration in the us-west-1 VPC. Create an Amazon Route 53 hosted zone. Create separate records for each ALEnable health checks to ensure high availability between Regions.

C. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPCreate an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPPlace the Auto Scaling group behind the ALB. Set up the same configuration in the us-west-1 VPCreate an Amazon Route 53 hosted zone. Create separate records for each ALB. Enable health checks and configure a failover routing policy for each record.

D. Create a VPC in us-east-1 and a VPC in us-west-1. Configure VPC peering. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in both VPCs. Create an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in both VPCs. Place the Auto Scaling group behind the ALB. Create an Amazon Route 53 hosted zone. Create a record for the ALB.

Suggested Answer: *C*

Community vote distribution

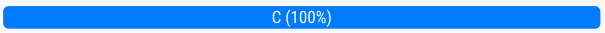

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct answer is C. Choice C meets the requirements for the application to be highly available and to dynamically scale to meet user traffic, as well as implementing a disaster recovery environment in the us-west-1 Region through active-passive failover.

In choice C, the company creates a VPC in us-east-1 and a VPC in us-west-1, and sets up an Application Load Balancer (ALB) and Auto Scaling group in both VPCs. The ALB extends across multiple Availability Zones in each VPC, and the Auto Scaling group deploys the EC2 instances across these Availability Zones. The Auto Scaling group is placed behind the ALB, which allows for automatic scaling of the instances to meet user traffic.

An Amazon Route 53 hosted zone is also created, with separate records for each ALB. Health checks are enabled for each record, and a failover routing policy is configured. This allows for active-passive failover between the two regions, ensuring high availability for the application.

upvoted 19 times

masetromain 2 years, 5 months ago

Choice A, B, and D do not fully meet the requirements of the disaster recovery environment in the us-west-1 Region and the failover routing policy because they do not include the necessary configurations for active-passive failover.

In choice A, the VPCs in us-east-1 and us-west-1 are peered and the Auto Scaling group and Application Load Balancer (ALB) are extended across multiple availability zones in both regions. However, there is no explicit failover routing policy configured, so it is not clear how the application would failover to the us-west-1 region in the event of an outage.

Choice B, the VPCs in us-east-1 and us-west-1 are separate, and the configuration is replicated in both regions but there is no explicit failover routing policy configured, so it is not clear how the application would failover to the us-west-1 region in the event of an outage.

upvoted 6 times

masetromain 2 years, 5 months ago

Choice D is similar to choice A, the VPCs in us-east-1 and us-west-1 are peered and the Auto Scaling group and Application Load Balancer (ALB) are extended across multiple availability zones in both regions. However, there is no explicit failover routing policy configured, so it is not clear how the application would failover to the us-west-1 region in the event of an outage.

Choice C is the correct answer as it includes all the necessary components for a disaster recovery environment in the us-west-1 region. It creates separate VPCs, Application Load Balancer, and Auto Scaling Group in both regions, and it enables health checks and configure a failover routing policy for each record. This ensures that in the event of an outage, the application can automatically failover to the us-west-1 region with minimal downtime.

upvoted 6 times

zozza2023 Highly Voted 2 years, 5 months ago

Selected Answer: C

active-passive failover==>a failover routing policy within route 53

upvoted 7 times

amministrazione Most Recent 10 months ago

C. Create a VPC in us-east-1 and a VPC in us-west-1. In the us-east-1 VPC, create an Application Load Balancer (ALB) that extends across multiple Availability Zones in that VPCreate an Auto Scaling group that deploys the EC2 instances across the multiple Availability Zones in the us-east-1 VPPlace the Auto Scaling group behind the ALB. Set up the same configuration in the us-west-1 VPCreate an Amazon Route 53 hosted zone. Create separate records for each ALB. Enable health checks and configure a failover routing policy for each record.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

It's C

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

C for DR

upvoted 2 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

Active-Passive failover with primary and secondary records in Route53
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html
https://d1tcczg8b21j1t.cloudfront.net/strapi-assets/32_Route_53_health_checks_4_64165fc533.png

upvoted 5 times

God_Is_Love 2 years, 3 months ago

VPC Peering is good for fully accessing all resources in a shared env but thats not asked here, so A and D gets eliminated. B does not mention the weighted routing config enable ment although setup is good. So answer is C

upvoted 3 times

zhangyu20000 2 years, 5 months ago

C is correct

upvoted 3 times

Question #70 *Topic 1*

A company has an environment that has a single AWS account. A solutions architect is reviewing the environment to recommend what the company could improve specifically in terms of access to the AWS Management Console. The company's IT support workers currently access the console for administrative tasks, authenticating with named IAM users that have been mapped to their job role.

The IT support workers no longer want to maintain both their Active Directory and IAM user accounts. They want to be able to access the console by using their existing Active Directory credentials. The solutions architect is using AWS IAM Identity Center (AWS Single Sign-On) to implement this functionality.

Which solution will meet these requirements MOST cost-effectively?

A. Create an organization in AWS Organizations. Turn on the IAM Identity Center feature in Organizations. Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Directory. Configure IAM Identity Center and set the AWS Managed Microsoft AD directory as the identity source. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.

B. Create an organization in AWS Organizations. Turn on the IAM Identity Center feature in Organizations. Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure IAM Identity Center and select the AD Connector as the identity source. Create permission sets and map them to the existing groups within the company's Active Directory.

C. Create an organization in AWS Organizations. Turn on all features for the organization. Create and configure a directory in AWS Directory Service for Microsoft Active Directory (AWS Managed Microsoft AD) with a two-way trust to the company's on-premises Active Directory. Configure IAM Identity Center and select the AWS Managed Microsoft AD directory as the identity source. Create permission sets and map them to the existing groups within the AWS Managed Microsoft AD directory.

D. Create an organization in AWS Organizations. Turn on all features for the organization. Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure IAM Identity Center and set the AD Connector as the identity source. Create permission sets and map them to the existing groups within the company's Active Directory.

Suggested Answer: *D*

Community vote distribution

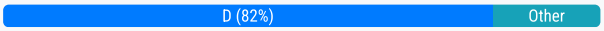

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: D

https://www.examtopics.com/discussions/amazon/view/69172-exam-aws-certified-solutions-architect-professional-topic-1/

You are correct, I apologize for the oversight. To meet the requirements of the IT support workers, option D would be the correct solution:

This option will first enable all features in AWS Organizations, then create and configure an AD Connector to connect to the company's on-premises Active Directory. Then, it will configure IAM Identity Center (AWS SSO) and set the AD Connector as the identity source, allowing the IT support workers to access the console using their existing Active Directory credentials. Finally, it will create permission sets and map them to the existing groups within the company's Active Directory. This solution will also be cost-effective as it does not involve creating a new directory in AWS Directory Service.

upvoted 22 times

dev112233xx Highly Voted 2 years, 3 months ago

Selected Answer: D

D is the correct answer.. B is wrong answer

From aws documentation:
Q: Which AWS accounts can I connect to IAM Identity Center?

You can add any AWS account managed using AWS Organizations to IAM Identity Center. You need to enable all features in your organizations to manage your accounts single sign-on.

upvoted 17 times

carpa_jo 1 year, 6 months ago

Source: https://aws.amazon.com/iam/identity-center/faqs/#product-faqs#iam-identity-center-faqs#identity-sources-and-applications-support
upvoted 1 times

scaag92 Most Recent 2 months, 3 weeks ago

Selected Answer: B

Opción B dice: "...Enciende la funcionalidad de IAM Identity Center en Organizations..." (Activa solo lo necesario para SSO).

Opción D dice: "...Enciende todas las funcionalidades para la organización..." (Activa todo lo que Organizations puede hacer).
Ambas B y D eligen la forma económica de conectar con AD (usando AD Connector). Pero la Opción B es más precisa y sigue mejor las buenas prácticas al activar solo la funcionalidad específica de IAM Identity Center que se necesita, en lugar de activar "todas" las funcionalidades de Organizations (como dice la D) que no son requeridas para esta tarea en particular. Por eso, B es la respuesta ligeramente mejor y más correcta.
upvoted 1 times

BelloMio 2 months, 4 weeks ago

Selected Answer: D

Answer is D:
I quote:
AWS Organizations is recommended, but not required, for use with IAM Identity Center. If you haven't set up an organization, you don't have to. If you've already set up AWS Organizations and are going to add IAM Identity Center to your organization, make sure that all AWS Organizations features are enabled
https://docs.aws.amazon.com/singlesignon/latest/userguide/identity-center-prerequisites.html
upvoted 1 times

hhiguita 3 months, 1 week ago

Selected Answer: B

All features is not required.
upvoted 1 times

BelloMio 2 months, 4 weeks ago

AWS Organizations is recommended, but not required, for use with IAM Identity Center. If you haven't set up an organization, you don't have to. If you've already set up AWS Organizations and are going to add IAM Identity Center to your organization, make sure that all AWS Organizations features are enabled

Quoting from:
https://docs.aws.amazon.com/singlesignon/latest/userguide/identity-center-prerequisites.html
upvoted 1 times

29fb203 3 months, 3 weeks ago

Selected Answer: B

All features is not required.
upvoted 1 times

LeoSantos121212121212121 3 months, 3 weeks ago

Selected Answer: B

Lower cost compared to AWS Managed Microsoft AD, since AD Connector does not require an additional managed directory service. Also, is answer D AWS Organizations does not require "all features" for IAM Identity Center to work with AD Connector. "All features" is needed for SCPs and governance, not for SSO setup.
upvoted 1 times

shmoeee 5 months ago

Selected Answer: D

"Need to turn on all features" didn't sound cost effective...but apparently it's a requirement to provide SSO
upvoted 1 times

JOJO9 6 months, 2 weeks ago

Selected Answer: B

Question asks "MOST cost-effectively". Turning on all features is free of charge but used resources will make up a cost.

D is wrong because: enabling All Features in AWS Organizations introduces governance tools that the company does not require, making it less cost-effective than B.
upvoted 3 times

amministrazione 10 months ago

D. Create an organization in AWS Organizations. Turn on all features for the organization. Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure IAM Identity Center and set the AD Connector as the identity source. Create permission sets and map them to the existing groups within the company's Active Directory.

upvoted 1 times

pk0619 6 months, 1 week ago

You need to enable all features for the organization to set up single sign-on for accounts.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B, Turn on the IAM Identity Center feature in Organizations... similar to D, but without enabling directy the SSO, you can't configure it...

upvoted 2 times

helloworldabc 9 months, 4 weeks ago

just D

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: D

Option D is the best because AWS FAQs asked the following question and answered: "Which AWS accounts can I connect to IAM Identity Center? You can add any AWS account managed using AWS Organizations to IAM Identity Center. You need to enable all features in your organizations to manage your accounts single sign-on." Link: https://aws.amazon.com/iam/identity-center/faqs/#product-faqs#iam-identity-center-faqs#identity-sources-and-applications-support.

With the clarification that enabling all features in AWS Organizations is necessary for integrating with IAM Identity Center, Option D becomes the most accurate and compliant solution. It correctly combines the need to enable all features in AWS Organizations with the use of an AD Connector for a direct connection to the company's on-premises Active Directory, which remains the most cost-effective way to leverage existing Active Directory credentials for AWS console access.

upvoted 1 times

LazyAutonomy 1 year, 4 months ago

Selected Answer: D

Most cost effective is D.

But C is also technically a valid solution that meets all the other requirements. A two way trust means AD users in the on-premise AD can be added to AD groups in the AWS-managed AD.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: D

A = you do not turn on AWS IdC feature only in AWS Orgs. It is either Consolidation billing or All features

B = same as above

C = requirements is to login users based on-premise AD, for this there is no need to AWS Managed AD with a local domain/directory and 2-way trust. AD Connector is enough and cheaper

D = correct

upvoted 5 times

marszalekm 1 year, 5 months ago

I love such questions, while both B and D seems reasonable, I thinking more about B because of this https://docs.aws.amazon.com/singlesignon/latest/userguide/get-set-up-for-idc.html

upvoted 1 times

Russs99 1 year, 6 months ago

Selected Answer: B

you can absolutely use an AD Connector as the identity source for AWS IAM Identity Center without turning on all features in your AWS organization. In fact, it's the most cost-effective and recommended approach if you only need single sign-on functionality with your existing on-premises Active Directory

upvoted 3 times

holymancolin 1 year, 7 months ago

Selected Answer: D

https://docs.aws.amazon.com/singlesignon/latest/userguide/prereq-orgs.html

```If you've already set up AWS Organizations and are going to add IAM Identity Center to your organization, make sure that all AWS Organizations features are enabled. When you create an organization, enabling all features is the default.```

upvoted 4 times
```

Question #71 *Topic 1*

A video streaming company recently launched a mobile app for video sharing. The app uploads various files to an Amazon S3 bucket in the us-east-1 Region. The files range in size from 1 GB to 10 GB.

Users who access the app from Australia have experienced uploads that take long periods of time. Sometimes the files fail to completely upload for these users. A solutions architect must improve the app's performance for these uploads.

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

A. Enable S3 Transfer Acceleration on the S3 bucket. Configure the app to use the Transfer Acceleration endpoint for uploads.

B. Configure an S3 bucket in each Region to receive the uploads. Use S3 Cross-Region Replication to copy the files to the distribution S3 bucket.

C. Set up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region.

D. Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3.

E. Modify the app to add random prefixes to the files before uploading.

Suggested Answer: *AD*

Community vote distribution

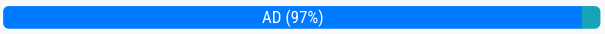

zozza2023 Highly Voted 2 years, 5 months ago

Selected Answer: AD

Transfer Accelerator + Multi-part uploads for files more 500MB

upvoted 12 times

OCHT Highly Voted 2 years, 2 months ago

Selected Answer: AD

Explanation for this .

B: Configuring an S3 bucket in each Region to receive the uploads and using S3 Cross-Region Replication to copy the files to the distribution S3 bucket may improve data durability and availability, but it does not address the issue of slow uploads from Australia.

C: Amazon Route 53 with latency-based routing can route the uploads to the nearest S3 bucket Region based on network latency, but it cannot guarantee faster upload speeds or better reliability.

E: Adding random prefixes to the files before uploading will not improve upload performance or reliability.

Thence, I select A and D.

upvoted 7 times

amministrazione Most Recent 10 months ago

A. Enable S3 Transfer Acceleration on the S3 bucket. Configure the app to use the Transfer Acceleration endpoint for uploads.
D. Configure the app to break the video files into chunks. Use a multipart upload to transfer files to Amazon S3.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: AD

A = correct (improve upload performance)
B = this could work along with C to improve performance, but this will not fix upload failure for files >5GB as you need multi-part upload
C = se answer B
D = correct (required to fix upload failures for >5GB files)
E = this could help with throttling which is not clearly stated as an issue

upvoted 1 times

chico2023 1 year, 10 months ago

Selected Answer: DE

Answer: A, D? Maybe. But I prefer D and E. Let me explain why:

Requirement is: "A solutions architect must improve the app's performance for these uploads."

Should we change S3 or the app? (or both?)

Depending on how you interpret this question, you might think on the app, then it should be D and E, seriously. And it DOES make sense. Bear with me here. If you break the files into chunks, you will still have to upload them, let's say 10GB. And here comes the option E, which helps improving uploads with PARALELLISM, and you didn't touch S3 to fix that, just the app :)

B and C would also work and would address the issue with users in Australia but it would change their design. I am not sure this is required, but in the real world, it's good to have options ;)

All in all, I personally would go with D, E, but AD and BC would also work.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: AD

its AD

upvoted 2 times

Maria2023 2 years ago

Selected Answer: AD

A and D satisfy the requirement

upvoted 1 times

SkyZeroZx 2 years, 1 month ago

Selected Answer: AD

Transfer Accelerator + Multi-part uploads for files more 500MB

Question similar to AWS Certified Solutions Architect Associate

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: AD

AD all day

upvoted 2 times

aqiao 2 years, 3 months ago

Selected Answer: AD

B is not suitable here, since it wants to improve upload experience, not download

upvoted 2 times

Musk 2 years, 5 months ago

I like AD but I am unsure. If the users in US don't complain about issues, it must be because multi-part upload is already enabled, otherwise it would fail 50% of the times. If only Australia users complain, it must be something else... Maybe A+B is a better option, although B is not the most cost efficient certainly.

upvoted 2 times

zhangyu20000 2 years, 5 months ago

AD is correct

upvoted 1 times

masetromain 2 years, 5 months ago

Selected Answer: AD

https://www.examtopics.com/discussions/amazon/view/74177-exam-aws-certified-solutions-architect-professional-topic-1/

The correct answers would be A and D.

A. Enabling S3 Transfer Acceleration on the S3 bucket and configuring the app to use the Transfer Acceleration endpoint for uploads will improve the app's performance for users in Australia by providing a fast and secure way to transfer large files over the Internet.

D. Configuring the app to break the video files into chunks and using a multipart upload to transfer files to Amazon S3, will improve the app's

performance for users in Australia by allowing them to upload large files in parallel, which can increase upload speed and reduce the risk of upload failures.

upvoted 4 times

masetromain 2 years, 5 months ago

B. Configuring an S3 bucket in each Region to receive the uploads and using S3 Cross-Region Replication to copy the files to the distribution S3 bucket is not the most cost-effective solution for this specific use case.

C. Setting up Amazon Route 53 with latency-based routing to route the uploads to the nearest S3 bucket Region is not a solution that would improve the performance of the uploads specifically for users in Australia.

E. Modifying the app to add random prefixes to the files before uploading will not improve the app's performance for users in Australia.

upvoted 1 times

hobokabobo 2 years, 3 months ago

yes, it will. Other options are more important, but sure random (rsp. any hash that distributes well) prefixes improve performance a lot.

upvoted 2 times

Question #72 *Topic 1*

An application is using an Amazon RDS for MySQL Multi-AZ DB instance in the us-east-1 Region. After a failover test, the application lost the connections to the database and could not re-establish the connections. After a restart of the application, the application re-established the connections.

A solutions architect must implement a solution so that the application can re-establish connections to the database without requiring a restart.

Which solution will meet these requirements?

A. Create an Amazon Aurora MySQL Serverless v1 DB instance. Migrate the RDS DB instance to the Aurora Serverless v1 DB instance. Update the connection settings in the application to point to the Aurora reader endpoint.

B. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.

C. Create a two-node Amazon Aurora MySQL DB cluster. Migrate the RDS DB instance to the Aurora DB cluster. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.

D. Create an Amazon S3 bucket. Export the database to Amazon S3 by using AWS Database Migration Service (AWS DMS). Configure Amazon Athena to use the S3 bucket as a data store. Install the latest Open Database Connectivity (ODBC) driver for the application. Update the connection settings in the application to point to the Athena endpoint

Suggested Answer: *B*

Community vote distribution

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: B

Amazon RDS Proxy is a fully managed database proxy service for Amazon Relational Database Service (RDS) that makes applications more scalable, resilient, and secure. It allows applications to pool and share connections to an RDS database, which can help reduce database connection overhead, improve scalability, and provide automatic failover and high availability.

upvoted 10 times

zhangyu20000 Highly Voted 2 years, 5 months ago

B is correct.

C: Aurora is useless, Proxy is pointing to existing RDS

upvoted 7 times

amministrazione Most Recent 10 months ago

B. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.

upvoted 1 times

pangchn 1 year, 2 months ago

Selected Answer: B

C is wrong since RDS proxy for Aurora cluster only support reader endpoint, where in question it doesn't mention the read-only as requirement

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

A = using Aurora MySQL Serverless will not fix the issue, also serverless V1 is not great with HA. If you are running a single instance (no read replicas) it will attempt to create a new DB Instance in the same AZ

B = correct (RDS Proxy in addition to pooling connections, makes applications more resilient to database failures by automatically connecting to a standby DB instance while preserving application connections and detects failover and routes requests to standby instance up to 66% faster failover time)

C = Creating and migrating to Aurora cluster is not needed, RDS Proxy is enough

D = this requires a lot of work

upvoted 5 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

it's a B

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

keyword = RDS proxy

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

Create an RDS proxy.

upvoted 1 times

klog 2 years, 4 months ago

Selected Answer: B

proxy will be a buffer

upvoted 1 times

masetromain 2 years, 5 months ago

Selected Answer: B

The correct solution is B. Create an RDS proxy. Configure the existing RDS endpoint as a target. Update the connection settings in the application to point to the RDS proxy endpoint.

An RDS proxy is a service that allows you to pool and share connections to an RDS database. By using an RDS proxy, your application can automatically reconnect to the database after a failover event, without the need to restart the application.

Solution A, migrating to Aurora Serverless, may not solve the problem because Aurora Serverless does not support Multi-AZ.
Solution C and D are not the correct solutions because it does not solve the problem of reconnecting to the database after a failover event.

upvoted 4 times

God_Is_Love 2 years, 3 months ago

What?? Aurora does not support Multi AZ ? its a blunder !

upvoted 5 times

chikorita 2 years ago

was about to point this

upvoted 1 times

BabaP 2 years ago

they are copying the answers from chatgpt

upvoted 7 times

k8s_Seoul 1 year, 9 months ago

masetromain ~> X

GPTromain ~> O lol

upvoted 1 times

SeemaDataReader 1 year, 6 months ago

Even if the person is copying from chatgpt, they are saving your time and giving some pointers.

upvoted 1 times

Question #73 *Topic 1*

A company is building a solution in the AWS Cloud. Thousands or devices will connect to the solution and send data. Each device needs to be able to send and receive data in real time over the MQTT protocol. Each device must authenticate by using a unique X.509 certificate.

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

A. Set up AWS IoT Core. For each device, create a corresponding Amazon MQ queue and provision a certificate. Connect each device to Amazon MQ.

B. Create a Network Load Balancer (NLB) and configure it with an AWS Lambda authorizer. Run an MQTT broker on Amazon EC2 instances in an Auto Scaling group. Set the Auto Scaling group as the target for the NLConnect each device to the NLB.

C. Set up AWS IoT Core. For each device, create a corresponding AWS IoT thing and provision a certificate. Connect each device to AWS IoT Core.

D. Set up an Amazon API Gateway HTTP API and a Network Load Balancer (NLB). Create integration between API Gateway and the NLB. Configure a mutual TLS certificate authorizer on the HTTP API. Run an MQTT broker on an Amazon EC2 instance that the NLB targets. Connect each device to the NLB.

Suggested Answer: *D*

Community vote distribution

C (97%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct solution is C. Set up AWS IoT Core. For each device, create a corresponding AWS IoT thing and provision a certificate. Connect each device to AWS IoT Core.

AWS IoT Core is a fully managed service that enables secure, bi-directional communication between internet-connected devices and the AWS Cloud. It supports the MQTT protocol and includes built-in device authentication and access control. By using AWS IoT Core, the company can easily provision and manage the X.509 certificates for each device, and connect the devices to the service with minimal operational overhead.

upvoted 23 times

masetromain 2 years, 5 months ago

Option A, setting up Amazon MQ queues and connecting each device to a queue, would require significant operational overhead to manage the queues and ensure that each device is properly authenticated and connected.

Option B and D, using a Network Load Balancer (NLB) with a Lambda authorizer or an Amazon API Gateway HTTP API with a mutual TLS certificate authorizer and running an MQTT broker on EC2 instances, would also introduce more operational complexity and overhead compared to using AWS IoT Core.

upvoted 6 times

MAZIADI Most Recent 10 months, 3 weeks ago

Selected Answer: C

AWS IoT Core: This service is specifically designed for managing IoT devices and supports the MQTT protocol natively. It provides built-in support for device authentication using X.509 certificates.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, use IoT Core

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: C

Option C is the most suitable solution as AWS IoT Core is specifically designed for IoT scenarios, including device management and secure communication. AWS IoT Core natively supports MQTT, a lightweight communication protocol ideal for IoT devices. It allows devices to connect securely with an individual X.509 certificate for authentication, significantly reducing operational overhead compared to managing a custom MQTT broker or other intermediate services. AWS IoT Core also simplifies device management and scaling, making it the best choice for the described use case.

upvoted 1 times

bjexamprep 1 year, 5 months ago

Selected Answer: C

I don't like C, but C might be the preferred answer.

There are thousands of devices. If C is the real answer, there should be a way to automatically create IOT thing and provision certificate. The answer seems implying to create IOT thing and provision certificates manually. If IoT core doesn't have this automation feature, this definitely is not the right answer in real life.

If there is this automation way and the question designer is expecting the exam taker to know this detail, that might be too specific for the exam takers.

D is ugly, and usually is not a correct answer in most question designs. But it provides a feasible way in the real life comparing with C.

upvoted 3 times

waoo 1 year, 11 months ago

□□□C

https://aws.amazon.com/cn/iot-core/faqs/?nc=sn&loc=5&dn=2

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

it's C

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

I choose C

upvoted 1 times

zejou1 2 years, 3 months ago

Selected Answer: C

https://docs.aws.amazon.com/iot/latest/developerguide/attach-to-cert.html

It is C, - you have to do this through IOT core, for the devices you need an AWS IOT "thing" and then provision a certificate for the thing. from there connect the device.

upvoted 2 times

forceli 2 years, 3 months ago

Selected Answer: A

-The AWS IoT Device SDKs support device communications using the MQTT

-Device connections to AWS IoT use X.509 client certificates

https://docs.aws.amazon.com/iot/latest/developerguide/iot-connect-devices.html

upvoted 1 times

forceli 2 years, 3 months ago

Sorry I meant "C"

upvoted 2 times

zozza2023 2 years, 5 months ago

Selected Answer: C

C is correct (less op overhead than A)

upvoted 2 times

zhangyu20000 2 years, 5 months ago

C is correct

upvoted 3 times

Question #74 *Topic 1*

A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.

What should the solutions architect do to create the solution?

A. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket. Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormation. Use AWS CloudFormation templates to provision resources.

B. Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormation. Use AWS CloudFormation templates to create stacks with approved resources.

C. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role. Assign the IAM service role to AWS CloudFormation during stack creation.

D. Provision resources in AWS CloudFormation stacks. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.

Suggested Answer: *B*

Community vote distribution

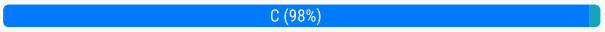

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: C

Tricky one. Question has a hint -"to enforce the new restriction on the IAM role" (note its not IAM policy as mentioned in option B) Creating a policy with approved resources first and assuming/applying that role to engineers will enforce. So C is correct. (B lacks enforcement, B is incorrect)

upvoted 18 times

rbm2023 Highly Voted 2 years, 1 month ago

Selected Answer: C

C is correct not B , AWS CloudFormation makes calls to create, modify, and delete those resources on their behalf. To separate permissions between a user and the AWS CloudFormation service, use a service role. AWS CloudFormation uses the service role's policy to make calls instead of the user's policy. For more information, see AWS CloudFormation service role . check this out .

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

Option B would allow engineers to provision resources using other methods outside of CloudFormation, which would not comply with the new company policy. This would make it difficult to enforce the new restriction on the IAM role that the engineers use for access.

upvoted 11 times

amministrazione Most Recent 10 months ago

C. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role. Assign the IAM service role to AWS CloudFormation during stack creation.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, use the IAM service role to execute the stack

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: C

Option C is the most effective solution. It involves updating the engineers' IAM role to only allow actions related to AWS CloudFormation, effectively preventing direct provisioning or modification of AWS resources outside of CloudFormation. By creating a service role (with permissions to provision approved resources) that CloudFormation assumes when executing templates, you enforce the provisioning of only approved resources through CloudFormation. This setup provides a clear separation of permissions: engineers can manage CloudFormation stacks but cannot directly create resources unless defined in a CloudFormation template and permitted by the service role.

Option B suggests updating the IAM policy to allow only the provisioning of approved resources and CloudFormation actions. This approach could theoretically work by explicitly listing allowed actions for specific AWS services in the IAM policy. However, it might be challenging to maintain and could inadvertently permit actions outside of CloudFormation, depending on the policy's specificity.

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

A = doesn't prevent to have a CloudFromation template with non-approved resources deployed

B = this doesn't prevent engineers to provision resources from console or cli

C = correct

D = doesn't prevent to provision non-approved resources or to provision only via CloudFormation

upvoted 2 times

subbupro 1 year, 6 months ago

B would be created generally in organization. C is fine , but more restriction , the user can only use the cloud formation stack sets only which is not good for organization level.

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: C

with B engineer will be able to directly provision resources without using of CF

upvoted 2 times

venvig 1 year, 10 months ago

Selected Answer: C

The two contenders are Option B and C.

Option B would allow the users to provision the approved resources without using CloudFormation (as the Users' IAM role would permission that). So, this violates the requirement.

Option C would ensure that Only Cloudformation can provision the resources. So, that's the correct answer.

upvoted 1 times

CuteRunRun 1 year, 10 months ago

Selected Answer: C

I prefer C, because you need to give permission to cloud formation

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C no doubt

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

C. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions.

upvoted 2 times

c73bf38 2 years, 4 months ago

Selected Answer: C

C IAM policy is allowing to provision of approved resources.

upvoted 3 times

Musk 2 years, 5 months ago

Selected Answer: C

B does not enfore CF, otherwise it would work.

upvoted 3 times

Untamables 2 years, 5 months ago

Selected Answer: C

C

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-control-access

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

upvoted 3 times

Nicocacik 2 years, 5 months ago

Selected Answer: C

You have to use a service role
upvoted 4 times

masetromain 2 years, 5 months ago

C. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role. Assign the IAM service role to AWS CloudFormation during stack creation.

This option is also correct, it is a way to restrict the access of engineers to only be able to perform AWS CloudFormation actions and provision only approved resources. By giving only permissions to the IAM role used by engineers for CloudFormation and creating a separate IAM role with permissions to provision approved resources and then assigning that role to CloudFormation during stack creation, we ensure that engineers can only provision the approved resources using CloudFormation.
upvoted 2 times

masetromain 2 years, 5 months ago

Both options B and C are correct.

Option B: Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormation. Use AWS CloudFormation templates to create stacks with approved resources.

Option C: Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role. Assign the IAM service role to AWS CloudFormation during stack creation.

Both options will enforce the new restriction on the IAM role that the engineers use for access, by limiting their access only to approved resources and only allowing them to provision resources using AWS CloudFormation. The specif
upvoted 1 times

Japanese1 1 year, 7 months ago

B works but is inappropriate.
You fail to consider that you NEED to use CFn for resource provisioning.
Option B does not meet the requirement to limit this.
upvoted 1 times

Question #75 *Topic 1*

A solutions architect is designing the data storage and retrieval architecture for a new application that a company will be launching soon. The application is designed to ingest millions of small records per minute from devices all around the world. Each record is less than 4 KB in size and needs to be stored in a durable location where it can be retrieved with low latency. The data is ephemeral and the company is required to store the data for 120 days only, after which the data can be deleted.

The solutions architect calculates that, during the course of a year, the storage requirements would be about 10-15 TB.

Which storage strategy is the MOST cost-effective and meets the design requirements?

A. Design the application to store each incoming record as a single .csv file in an Amazon S3 bucket to allow for indexed retrieval. Configure a lifecycle policy to delete data older than 120 days.

B. Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale. Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.

C. Design the application to store each incoming record in a single table in an Amazon RDS MySQL database. Run a nightly cron job that runs a query to delete any records older than 120 days.

D. Design the application to batch incoming records before writing them to an Amazon S3 bucket. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data. Configure a lifecycle policy to delete the data after 120 days.

Suggested Answer: *B*

Community vote distribution

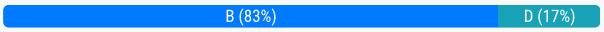

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The most cost-effective and efficient solution that meets the design requirements would be option B, Design the application to store each incoming record in an Amazon DynamoDB table properly configured for the scale. Configure the DynamoDB Time to Live (TTL) feature to delete records older than 120 days.

DynamoDB is a NoSQL key-value store designed for high scale and performance. It is fully managed by AWS and can easily handle millions of small records per minute. Additionally, with the TTL feature, you can set an expiration time for each record, so that the data can be automatically deleted after the specified time period.

upvoted 23 times

masetromain 2 years, 5 months ago

Option A, storing each incoming record as a single .csv file in an Amazon S3 bucket, would not be a good option because it would be difficult to retrieve individual records from the .csv files, and will likely increase the cost of data retrieval.

Option C, storing each incoming record in a single table in an Amazon RDS MySQL database, would be a more expensive option as RDS is typically more expensive than DynamoDB. Additionally, running a cron job to delete old data could lead to additional operational overhead.

Option D, storing incoming records in batches in an S3 bucket, would be a less efficient option as it would require additional processing and parsing of the data to retrieve individual records.

upvoted 7 times

dkx Highly Voted 1 year, 11 months ago

A. No, because millions of writes to a single .csv file would cause read and write latency

B. Yes, because DynamoDB can support peaks of more than 20 million requests per second.

C. No, because creating nightly cron is unnecessary, and a relation database isn't designed to ingest millions of small records per minute

D. No, because S3 supports 210,000 PUT requests per minute (3,500 requests per second * 60 seconds per min) which is far less than 1,000,000+ writes per minute

upvoted 6 times

ahhatem 6 months, 2 weeks ago

Actually, the limit you mentioned for point D is per prefix or path…. Not the whole bucket. With proper data distribution across prefixes it can accommodate easily for the load mentioned.

upvoted 2 times

jimee11 Most Recent 1 month, 3 weeks ago

Selected Answer: B

Read the requirements: MOST cost-effective and meets the design requirements. Note "retrieved with low latency". DynamoDB latency is single digits, where as S3 is 100-200 milliseconds.

upvoted 1 times

vmia159 3 months, 2 weeks ago

Selected Answer: D

For those who said B, how many WCU is needed for dynamoDB?

Given:

1 million records per minute

4KB per record

This translates to approximately 16,667 records per second (1,000,000 / 60)

For DynamoDB WCU calculation:

1 WCU = 1 write per second for items up to 1KB

For items larger than 1KB, the WCU is rounded up to the next 1KB

For 4KB items, each write will consume 4 WCUs

Therefore:

WCUs needed = (Records per second) × (Item size in KB rounded up)

WCUs = 16,667 × 4

WCUs = 66,668 WCUs

First, you need to increase the quotas for that table by submitting a support ticket.

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html

Second, this is very expensive.

Obviously, combine it with kinesis data agent and firehorse that write to S3 will be much reliable options but it will increase the cost significantly. But still cheaper than the dynamo db options.

https://calculator.aws/#/estimate?id=87f1df21449660b0b9d61a6c1153632b1983d2e4

upvoted 1 times

soulation 3 months, 3 weeks ago

Selected Answer: D

Option B is too expensive.

upvoted 1 times

sergza 6 months, 2 weeks ago

Selected Answer: D

If you really think about being cost effective than Option D is the right choice

upvoted 1 times

Heman31in 6 months, 3 weeks ago

Selected Answer: D

Why Option D Might Be Cost-Effective:

Lower Storage Costs:

S3 storage is generally cheaper than DynamoDB when dealing with large amounts of data (e.g., $0.023/GB/month for S3 Standard vs. $0.25/GB/month for DynamoDB on-demand).

Batching Reduces API Call Costs:

By batching multiple records into a single object, you reduce the number of PUT requests to S3. This can lead to lower API costs compared to writing each record individually to DynamoDB.

Lifecycle Policies for Data Expiry:

S3 lifecycle policies automatically clean up data older than 120 days, similar to DynamoDB's TTL feature.

upvoted 1 times

amministrazione 10 months ago

D. Design the application to batch incoming records before writing them to an Amazon S3 bucket. Update the metadata for the object to contain the list of records in the batch and use the Amazon S3 metadata search feature to retrieve the data. Configure a lifecycle policy to delete the data after

120 days.
upvoted 1 times

ahhatem 1 year ago

Selected Answer: B

Obviously it is DynamoDB. Although as a side node I would say it is probably a very bad choice as it would be astronomically expensive for millions of writes per minute…. A Kinesis Data Streams would make much more sense especially that the data is only needed for 3 months…
upvoted 2 times

ahhatem 6 months, 2 weeks ago

After a second thought, I am not sure it is B. D would be much cheaper if it means that objects buffered and combined before write. But the word "batch" doesn't make me comfortable, batching means writing the objects in one go… nothing implies the objects would be combined …
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B, dynamodb is the best option
upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: B

For small records less than 4 KB, DynamoDB can efficiently handle the ingestion of millions o records per minute from devices around the world, meeting the application's design requirements for low-latency data access. Additionally, DynamoDB's Time to Live (TTL) feature allows for automatic deletion of items after a specific period, aligning with the requirement to store data for only 120 days.
upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

A = S3 is not great with small files and searching for data based on index (a common pattern is to store object metadata in a database like DDB, OpenSearch or RDS/Aurora). Many small files can lead to high costs for retrieval
B = correct
C = single-table design, high volume write/retrieval os small object and no need for complex query are better served and cost less with DDB rather than RDS
D = more efficient than A, but still S3 metadata search feature is limited
upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: B

see uC6rW1aB's answer
upvoted 1 times

vjp_training 1 year, 9 months ago

Selected Answer: B

B is the best for cost-effective.
D is more cost for S3 request
upvoted 1 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: B

Ref: https://aws.amazon.com/dynamodb/pricing/on-demand/
DynamoDB read requests can be either strongly consistent, eventually consistent, or transactional. A strongly consistent read request of up to 4 KB requires one read request unit. For items larger than 4 KB, additional read request units are required.
upvoted 3 times

uC6rW1aB 1 year, 9 months ago

for a US East write object price:
S3 Standard put object per thound cost $0.005 -> 1 million put cost $5 (per minutes in this situation)
Dynamo DB 1 million write cost $1.25 is a lot of cheaper
upvoted 4 times

Gmail78 1 year, 10 months ago

Selected Answer: D

Dinamo DB is at least 5X more expensive than S3 for this use case. There are million of writing and each is 4K, total disk space is 10-15TB.
upvoted 1 times

vn_thanhtung 1 year, 10 months ago

D - S3 metadata search feature does not exist

upvoted 2 times

Soweetadad 1 year, 10 months ago

Selected Answer: D

Although both B and D are correct, Option D is more cost effective.

upvoted 1 times

Question #76 *Topic 1*

A retail company is hosting an ecommerce website on AWS across multiple AWS Regions. The company wants the website to be operational at all times for online purchases. The website stores data in an Amazon RDS for MySQL DB instance.

Which solution will provide the HIGHEST availability for the database?

A. Configure automated backups on Amazon RDS. In the case of disruption, promote an automated backup to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.

B. Configure global tables and read replicas on Amazon RDS. Activate the cross-Region scope. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.

C. Configure global tables and automated backups on Amazon RDS. In the case of disruption, use AWS Lambda to copy the read replicas from one Region to another Region.

D. Configure read replicas on Amazon RDS. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.

Suggested Answer: *D*

Community vote distribution

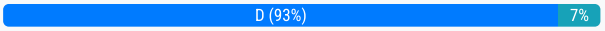

zejou1 Highly Voted 2 years, 3 months ago

Selected Answer: D

This really should be multi-az but you could move to it w/ D.
Here is the key to this one though; Highest Availability - the read replica is an asynchronous copy, while backup is a "time". Easier to do the read replica, and flip the switches than to reload from backup. Global Tables relate to DynomoDB https://disaster-recovery.workshop.aws/en/services/databases/dynamodb/dynamo-global-table.html
Little handy "DR" guide
upvoted 15 times

amministrazione Most Recent 10 months ago

D. Configure read replicas on Amazon RDS. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source.
upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: D

A = you cannot promote an automated backup to a standalone DB (you restore a backup into a new DB instance instead). Creating a read replica could help in this scenario in case it is cross-region. This is not specified
B = RDS does not support global table, copying a read replicas from a region to another make no sense to me
C = see B
D = correct
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

D for sure
upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: D

There is Aurora Global Database, DynamoDB Global Tables and the question is about RDS for MySQL DB Instance.
https://jayendrapatil.com/aws-aurora-global-database-vs-dynamodb-global-tables/
So, options B and C are not acceptable.
Option D refers to using a cross-region replication for disaster recovery which can be found here https://disaster-recovery.workshop.aws/en/services/databases/rds/rds-cross-region.html
Following article demonstrates a similar scenario using RDS for SQL Server
https://aws.amazon.com/blogs/database/use-cross-region-read-replicas-with-amazon-relational-database-service-for-sql-server/

The design seems to be what we are looking in terms of option D.
https://d2908q01vomqb2.cloudfront.net/887309d048beef83ad3eabf2a79a64a389ab1c9f/2022/11/15/dbblog-2614-image001.png
upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: D

D makes the most sense
upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: D

No global tables concept in RDS, B,C are eliminated. A is wrong in terms of backing up Db copy to a standalone instance ? D provides read replicas for reading and also swtiches as a failiover in times of disruption and becomes primary. this is how HA can be maintained. D is correct.
upvoted 3 times

spd 2 years, 4 months ago

Selected Answer: D

MySQL - Read Replica. In this case, this is not aurora so not the global table option and hence can not be B and C
upvoted 2 times

sambb 2 years, 4 months ago

I haven't found any information about a "global table" for RDS.
Global tables are for DynamoDB. For Aurora, it's called "global databases".
RDS for MySQL supports cross-region read replicas https://aws.amazon.com/fr/blogs/aws/cross-region-read-replicas-for-amazon-rds-for-mysql/, so D has a better availability than A.
upvoted 2 times

icassp 2 years, 5 months ago

Selected Answer: D

for B,C, Amazon RDS does not support global tables yet. Only Aurora supports.
upvoted 4 times

AlanKrish 2 years, 4 months ago

Is Aurora not part of RDS? You can choose Aurora's compatibility with MySQL and PostreSQL).
upvoted 1 times

zhangyu20000 2 years, 5 months ago

D is correct
upvoted 3 times

masetromain 2 years, 5 months ago

https://www.examtopics.com/discussions/amazon/view/69438-exam-aws-certified-solutions-architect-professional-topic-1/
upvoted 1 times

masetromain 2 years, 5 months ago

It is possible that some people may think that option D. Configure read replicas on Amazon RDS. In the case of disruption, promote a cross-Region and read replica to be a standalone DB instance. Direct database traffic to the promoted DB instance. Create a replacement read replica that has the promoted DB instance as its source. is the best solution, as it also utilizes read replicas and cross-Region promotion to minimize downtime. However, it is important to consider that while this solution provides high availability, it doesn't provide the same level of automatic replication that global tables do. In case of a disruption, there is a risk of data loss during the manual switchover.
and also with option D, you are still working with a single point of failure, the primary database, while in option B you have multiple copies of your data distributed across different regions, so in case of a failure you can switch over to one of the replicas without loss of data.
upvoted 2 times

Shahul75 2 years, 4 months ago

B is not right. Only Aurora has global tables. RDS don't
upvoted 1 times

[Removed] 2 years, 4 months ago

Cant be B due to global tables, ReadReplicas are supported with RDS and other options of restoring from backup do not create high availability
upvoted 1 times

masetromain 2 years, 5 months ago

Selected Answer: B

The correct answer is option B. Configuring global tables and read replicas on Amazon RDS with the cross-Region scope enabled provides the highest availability for the database. In case of disruption, the company can use AWS Lambda to copy the read replicas from one Region to another Region, ensuring that the website remains operational at all times. This solution provides automatic failover across multiple regions and allows for fast recovery in case of a disruption.

Option A involves promoting an automated backup to be a standalone DB instance and creating a replacement read replica that has the promoted DB instance as its source. This solution is less efficient since it requires manual intervention and additional steps to promote the backup and create a replacement read replica.
upvoted 2 times

Sarutobi 2 years, 4 months ago
If the disruption is an outage that takes the Region offline completely, how could we use Lambda to copy the read replica from the Region that is no longer available to the backup to another Region?
upvoted 1 times

masetromain 2 years, 5 months ago
Option C involves configuring global tables and automated backups on Amazon RDS. This solution is less efficient since it does not provide automatic failover across multiple regions and requires additional steps to copy the read replicas from one Region to another Region using AWS Lambda.

Option D involves configuring read replicas on Amazon RDS. In the case of disruption, promoting a cross-Region and read replica to be a standalone DB instance. This solution is less efficient than Option B since it does not provide automatic failover across multiple regions and requires manual intervention to promote the read replica to a standalone instance.
upvoted 1 times

bcx 2 years ago
In fact global tables is a Dynamo DB thing. And RDS has Aurora Global Database. In this case Aurora is out of the question, it says RDS MySql, not Aurora (RDS) MySQL.
upvoted 1 times

Question #77 *Topic 1*

Example Corp. has an on-premises data center and a VPC named VPC A in the Example Corp. AWS account. The on-premises network connects to VPC A through an AWS Site-To-Site VPN. The on-premises servers can properly access VPC A. Example Corp. just acquired AnyCompany, which has a VPC named VPC B. There is no IP address overlap among these networks. Example Corp. has peered VPC A and VPC B.

Example Corp. wants to connect from its on-premise servers to VPC B. Example Corp. has properly set up the network ACL and security groups.

Which solution will meet this requirement with the LEAST operational effort?

A. Create a transit gateway. Attach the Site-to-Site VPN, VPC A, and VPC B to the transit gateway. Update the transit gateway route tables for all networks to add IP range routes for all other networks.

B. Create a transit gateway. Create a Site-to-Site VPN connection between the on-premises network and VPC B, and connect the VPN connection to the transit gateway. Add a route to direct traffic to the peered VPCs, and add an authorization rule to give clients access to the VPCs A and B.

C. Update the route tables for the Site-to-Site VPN and both VPCs for all three networks. Configure BGP propagation for all three networks. Wait for up to 5 minutes for BGP propagation to finish.

D. Modify the Site-to-Site VPN's virtual private gateway definition to include VPC A and VPC B. Split the two routers of the virtual private getaway between the two VPCs.

Suggested Answer: *D*

Community vote distribution

A (91%) 6%

rbm2023 Highly Voted 2 years, 1 month ago

Selected Answer: A

https://docs.aws.amazon.com/pt_br/whitepapers/latest/aws-vpc-connectivity-options/aws-transit-gateway-vpn.html
Transit gateway is an AWS managed high availability and scalability regional network transit hub used to interconnect VPCs and customer networks. AWS Transit Gateway + VPN, using the Transit Gateway VPN Attachment, provides the option of creating an IPsec VPN connection between your remote network and the Transit Gateway over the internet, as shown in the following picture.
https://docs.aws.amazon.com/images/whitepapers/latest/aws-vpc-connectivity-options/images/image4.png
Option A is the correct answer since the transit gateway will allow both VPCs to connect to the on premises network.
Option B suggests the same feature but is using the Transit Gateway in a incorrect way. The soul purpose of the gateway is to have point for interconnectivity.
upvoted 10 times

Tunstim Highly Voted 2 years, 1 month ago

For those that have written SAP-C02, how relevant are these questions to the real exam questions? After adequate preparation, I wanted to truly test my knowledge before dabbling into the exam and would really appreciate anyone's candid opinion.
Thanks.
upvoted 5 times

chikorita 1 year, 9 months ago

please reply to him
upvoted 2 times

amministrazione Most Recent 10 months ago

A. Create a transit gateway. Attach the Site-to-Site VPN, VPC A, and VPC B to the transit gateway. Update the transit gateway route tables for all networks to add IP range routes for all other networks.
upvoted 1 times

jceleste1 10 months, 1 week ago

After all, what is the right answer A or D ?
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A, Transit Gateway

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: A

Option A is the most straightforward and effective solution. A transit gateway acts as a cloud router that simplifies network topology and connectivity between on-premises networks, VPCs, and other AWS services. By attaching both VPCs (A and B) and the Site-to-Site VPN to a single transit gateway and updating the route tables accordingly, Example Corp. can enable seamless communication between its on-premises network and both VPCs. This approach minimizes operational effort by centralizing network management and eliminating the need for complex routing configurations or multiple VPN connections.

Option D proposes modifying the Site-to-Site VPN's virtual private gateway to include both VPC A and VPC B. However, a virtual private gateway cannot be directly shared or split between VPCs in the manner described. This option misunderstands the architecture of AWS networking components and their capabilities.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: A

A = correct

B = if you setup a second VPN you do not need a TGW

C = peering does not allow edge-to-edge routing (aka VPC B cannot access on-premise via VPC A and vice versa)

D = Virtual Private Gateway is specific to a single VPC

upvoted 2 times

Russs99 1 year, 10 months ago

Selected Answer: A

reluctantly selecting option A. these answers do not take into consideration that the On-promises already has a peered connection to VPC A through the existing site to site

upvoted 2 times

CuteRunRun 1 year, 10 months ago

Selected Answer: A

I think A is right, I do not know why other guys select D

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

surely A

upvoted 1 times

Parsons 2 years, 2 months ago

Selected Answer: A

A is the best option.

Creating a transit gateway and attaching Site-to-Site VPN, VPC A, and VPC B to the transit gateway would enable the on-premise servers to access VPC B with minimal operational effort. The transit gateway route tables would need to be updated with IP range routes for all the other networks to enable communication between the VPCs and the on-premises servers.

upvoted 2 times

Arnaud92 2 years, 3 months ago

Selected Answer: A

Solution A is the only one possible solution

upvoted 1 times

Arnaud92 2 years, 3 months ago

B is impossible : When you create a S2S VPN connection, it's between 2 entites (here, the onprem and VPC B). It says that they connect the onprem to VPCB with S2SVPN AND THEN to a TGW, it's not possible to connect a S2S VPN from onprem to VPC to a TGW (it"s a 3 entities). You can however connect a S2S VPN to a TGW (onprem to TGW) (which is solution A).

C : Does not work, there is no transitivity on AWS. S2S VPN cannot reach VPC B through VPC A

D is impossible : There is no magic, you cannot "split" router (that does not exist). VGW is attach to a single VPC. A S2S VPN cannot multiplex VPC

upvoted 4 times

Arnaud92 2 years, 3 months ago

A : the best (and the only one possible) answer : When you have 2 VPC, you have multiple solution to connect to onprem :

- Create 2 S2S VPN (1 for each VPC)

- or Create a TGW, attach both VPC to it and attach S2S VPN to it too
- or Create a third VPC (VPC routing), and peer VPC A with VPC routing, VPC B to VPC routing, attach a S2S VPN to VPC routing and use a NVA on VPC routing to route trafic. NVA can do transitivity.
Here, solution A is one of the possible answer
upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: A

A. Create a transit gateway. Attach the Site-to-Site VPN
upvoted 1 times

dev112233xx 2 years, 3 months ago

Selected Answer: A

A makes sense to me
upvoted 1 times

taer 2 years, 3 months ago

Selected Answer: A

A for me
upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: B

A has this wierd wording - attaching S-S VPN ? transit gateway attaches to VPCs only not S-S vpn. A is wrong. Since VPC A and VPC B are already peered, the easiest solution to connect from the on-premises servers to VPC B would be to create another Site-to-Site VPN connection between the on-premises data center and VPC B. This would require minimal operational effort, as the existing VPN connection with VPC A can remain unchanged.
upvoted 1 times

God_Is_Love 2 years, 3 months ago

oops this is wrong..VPN can be attached...
upvoted 1 times

God_Is_Love 2 years, 3 months ago

Moderator, please delete this comment..
upvoted 1 times

God_Is_Love 2 years, 3 months ago

https://docs.aws.amazon.com/vpn/latest/s2svpn/how_it_works.html
When you create a virtual private gateway, you can specify the private Autonomous System Number (ASN) for the Amazon side of the gateway. If you don't specify an ASN, the virtual private gateway is created with the default ASN (64512). You cannot change the ASN after you've created the virtual private gateway. Due to this reason, So A is not possible (with least effort). Answer should be B.
upvoted 1 times

Arnaud92 2 years, 3 months ago

THe VGW for VPCA is no more needed on A because you attach the VPCA to the TGW.
The ASN will be on the TGW attachment with the S2S VPN.
This is the best solution.
In the meantime, B is impossible. When you create a S2S VPN connection, it's between 2 entites (here, the onprem and VPC B). It says that they connect the onprem to VPCB with S2SVPN AND THEN to a TGW, it's not possible to connect a S2S VPN from onprem to VPC to a TGW. You can however connect a S2S VPN to a TGW (onprem to TGW).
upvoted 1 times

spd 2 years, 4 months ago

Selected Answer: A

TGW is the solutions
upvoted 1 times

Question #78 *Topic 1*

A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only.

The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits.

What should the company do to modify the application to send email messages from Amazon SES?

A. Configure the application to connect to Amazon SES by using TLS Wrapper. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Attach the IAM role to an Amazon EC2 instance.

B. Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.

C. Configure the application to use the SES API to send email messages. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Use the IAM role as a service role for Amazon SES.

D. Configure the application to use AWS SDKs to send email messages. Create an IAM user for Amazon SES. Generate API access keys. Use the access keys to authenticate with Amazon SES.

Suggested Answer: *A*

Community vote distribution

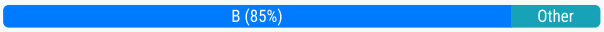

scuzzy2010 Highly Voted 2 years, 4 months ago

Selected Answer: B

B is correct.
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html
STARTTLS supports ports 25, 587, and 2587
TLSWRAPPER supports ports 465 and 2465
upvoted 19 times

God_Is_Love 2 years, 3 months ago

FYI Amazon SES supports STARTTLS encryption over port 587, which is the recommended port for email transmission. But existing port 25 can be configured too as in this case as the migration came from SMTP port 25
upvoted 5 times

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: B

In this scenario, you should use Amazon SES SMTP interface to send emails because the application can use SMTP only.
https://docs.aws.amazon.com/ses/latest/dg/send-email-smtp.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-credentials.html
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html
upvoted 9 times

amministrazione Most Recent 10 months ago

B. Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.
upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: B

Here's why option B is the correct choice:
STARTTLS Support: Amazon SES supports STARTTLS, a protocol command used to upgrade an existing insecure connection to a secure connection using TLS (Transport Layer Security). This is crucial since the legacy SMTP server does not support TLS, and STARTTLS can be used to initiate a secure connection.

SMTP Credentials: Amazon SES requires authentication to send emails through its SMTP interface. This is achieved by using SMTP credentials, which are different from AWS access keys. SMTP credentials can be obtained from the Amazon SES console and are used to authenticate with the Amazon SES SMTP endpoint.
Operational Simplicity: This approach allows the application to continue using SMTP for sending emails, which aligns with the application's existing capabilities. By using STARTTLS, the application can upgrade its connection to Amazon SES to a secure one, ensuring compliance with security best practices without significant changes to the application's email sending functionality.
upvoted 2 times

LazyAutonomy 1 year, 4 months ago

Selected Answer: A

Terrible Q. All answers are wrong.
A is wrong because you cannot send emails through SES SMTP using SMTP credentials derived from temporary STS tokens (ie IAM roles). Must use an IAM user access keys to derive creds.
B is wrong because the question imposes a constraint that prevents us from selecting an answer that requires upgrading or modifying the application itself. Could you just offload SMTP STARTTLS/AUTH to the local sendmail/postfix daemon? Maybe, if it were Linux, but what if it's Windows? Cygwin? WSL?
C & D - wrong, for a similar rationale as B.

But the question designer OBVIOUSLY doesn't know that IAM roles can't be used for SES SMTP auth, because these questions are written by inexperienced, unqualified people who are not themselves architects or engineers.
upvoted 2 times

> **LazyAutonomy** 1 year, 4 months ago
> To be fair, the question says this:
>
> "The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only."
>
> The question doesn't say the application cannot handle STARTTLS or SMTP AUTH. In theory, if an application claims to support SMTP, then it should support all features of SMTP, which includes STARTTLS and AUTH. It only says the legacy SMTP server cannot handle TLS. So I suppose perhaps B is correct after all :-)
> upvoted 3 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

A = this sends email via SES API while application can use SMTP only
B = correct
C = this sends email via SES API while application can use SMTP only
D = this sends email via SES SDK (API) while application can use SMTP only
upvoted 2 times

> **ninomfr64** 1 year, 5 months ago
> Need to correct my comment on A. This is a TLS Wrapper (A) vs STARTTLS (B), where STARTTLS allows initiating an encrypted connection by first establishing an unencrypted connection. While TLS Wrapper is a means of initiating an encrypted connection without first establishing an unencrypted connection (it's the client's responsibility to connect to the endpoint using TLS, and to continue using TLS for the entire conversation). As our app con only work with SMTP we should go for B
> upvoted 2 times

edder 1 year, 7 months ago

Selected Answer: B

The correct answer is B.
A: We are unable to obtain authentication information.
C,D: Does not meet SMTP requirements.

B: This is the correct procedure.

https://repost.aws/knowledge-center/ses-set-up-connect-smtp
https://docs.aws.amazon.com/ses/latest/dg/security-protocols.html
upvoted 1 times

totten 1 year, 9 months ago

Selected Answer: B

Here's why option B is the correct choice:

SMTP Protocol: The legacy SMTP server uses the SMTP protocol, and Amazon SES provides an SMTP interface for sending email, which is suitable for your application.

STARTTLS: Using STARTTLS ensures that your communication with Amazon SES is encrypted, which is a best practice for secure email transmission.

SMTP Credentials: Amazon SES SMTP credentials are required to authenticate your application with Amazon SES when sending emails. These credentials include an SMTP username and password.
upvoted 2 times

totten 1 year, 9 months ago
Option A mentions TLS Wrapper, which isn't a standard approach when using Amazon SES for sending email. Amazon SES supports STARTTLS for secure communication.

Option C suggests using the SES API, which is a valid approach but requires code modifications to use the SES API instead of SMTP. Since your application can only use SMTP, this option might involve significant code changes.

Option D mentions using AWS SDKs and IAM users, which is more suitable for programmatic access to SES but not for legacy SMTP applications that can only send via SMTP.

Therefore, Option B is the most appropriate choice for configuring your application to send email messages from Amazon SES while preserving the SMTP protocol and ensuring secure communication.
upvoted 4 times

CuteRunRun 1 year, 10 months ago
Selected Answer: A
I selecte A
upvoted 1 times

NikkyDicky 1 year, 12 months ago
Selected Answer: B
It's B - to preserve SMTP protocol
upvoted 1 times

SkyZeroZx 2 years ago
Selected Answer: B
B because is "legacy" app then use properties to set SMTP
keyword === Obtain Amazon SES SMTP credentials
upvoted 1 times

F_Eldin 2 years, 1 month ago
Selected Answer: A
https://aws.amazon.com/blogs/big-data/query-and-visualize-aws-cost-and-usage-data-using-amazon-athena-and-amazon-quicksight/
upvoted 1 times

rbm2023 2 years, 1 month ago
Selected Answer: B
Option A states that the company would require to do more changes in the application than a replatform migration strategy where we are supposed to migrate the application with minimal changes. In Option A using the TLS wrapper would require an additional layer of software (stunnel) to be installed and configured on the EC2 instance, which may introduce additional complexity and management overhead.
In option B, we need to configure the application to connect to SES using STARTLS using SMTP credentials, since the legacy SMTP server does not support TLS encryption. This would require minimal change to the application.
upvoted 3 times

Cassa 2 years, 2 months ago
Selected Answer: B
To set up a STARTTLS connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 25, 587, or 2587, issues an EHLO command, and waits for the server to announce that it supports the STARTTLS SMTP extension. The client then issues the STARTTLS command, initiating TLS negotiation. When negotiation is complete, the client issues an EHLO command over the new encrypted connection, and the SMTP session proceeds normally

To set up a TLS Wrapper connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 465 or 2465. The server presents its certificate, the client issues an EHLO command, and the SMTP session proceeds normally.

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: B

B. Configure the application to connect to Amazon SES by using STARTTLS.

upvoted 1 times

Dimidrol 2 years, 3 months ago

Selected Answer: B

B , https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html

upvoted 3 times

dev112233xx 2 years, 3 months ago

Selected Answer: A

B is wrong because STARTTLS uses port 25 and EC2 instances can't send outbound traffic through port 25 (you must ask AWS to allow port 25)

upvoted 2 times

F_Eldin 2 years ago

https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html

says:

"Amazon Elastic Compute Cloud (Amazon EC2) throttles email traffic over port 25 by default. To avoid timeouts when sending email through the SMTP endpoint from EC2,submit a Request to Remove Email Sending Limitations"

And the question explicitly says:

"The company has lifted the SES limits."

upvoted 2 times

Question #79 *Topic 1*

A company recently acquired several other companies. Each company has a separate AWS account with a different billing and reporting method. The acquiring company has consolidated all the accounts into one organization in AWS Organizations. However, the acquiring company has found it difficult to generate a cost report that contains meaningful groups for all the teams.

The acquiring company's finance team needs a solution to report on costs for all the companies through a self-managed application.

Which solution will meet these requirements?

A. Create an AWS Cost and Usage Report for the organization. Define tags and cost categories in the report. Create a table in Amazon Athena. Create an Amazon QuickSight dataset based on the Athena table. Share the dataset with the finance team.

B. Create an AWS Cost and Usage Report for the organization. Define tags and cost categories in the report. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

C. Create an Amazon QuickSight dataset that receives spending information from the AWS Price List Query API. Share the dataset with the finance team.

D. Use the AWS Price List Query API to collect account spending information. Create a specialized template in AWS Cost Explorer that the finance department will use to build reports.

Suggested Answer: *D*

Community vote distribution

A (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

The correct solution is A.

Creating an AWS Cost and Usage Report for the organization and defining tags and cost categories in the report will allow for detailed cost reporting for the different companies that have been consolidated into one organization. By creating a table in Amazon Athena and an Amazon QuickSight dataset based on the Athena table, the finance team will be able to easily query and generate reports on the costs for all the companies. The dataset can then be shared with the finance team for them to use for their reporting needs.

Option B is not correct because it does not provide a way to query and generate reports on the costs for all the companies.

Option C is not correct because it only provides spending information from the AWS Price List Query API and does not provide detailed cost reporting for the different companies.

Option D is not correct because it only uses the AWS Price List Query API and does not provide a way to query and generate reports on the costs for all the companies.

upvoted 14 times

moota Highly Voted 2 years, 4 months ago

Selected Answer: A

I can customize reporting in Cost Explorer but cannot find how to do templates.

upvoted 7 times

amministrazione Most Recent 10 months ago

A. Create an AWS Cost and Usage Report for the organization. Define tags and cost categories in the report. Create a table in Amazon Athena. Create an Amazon QuickSight dataset based on the Athena table. Share the dataset with the finance team.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: A

A = correct

B = there isn't specialized templates in AWS Cost Explorer. It provides default reports, but also enables you to change the filters and constraints used to create the reports. You can save the reports that you made as a bookmark, download the CSV file, or save them as a report

C & D = AWS Price List provides a catalog of the products and prices for AWS services that you can purchase on AWS, not cost of your resources

upvoted 3 times

CuteRunRun 1 year, 10 months ago

Selected Answer: A

I prefer A

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

its n A

upvoted 1 times

Maria2023 2 years ago

Selected Answer: A

I vote A mostly because there is no template option in Cost Explorer and A is the only other option which covers the scenario

upvoted 2 times

F_Eldin 2 years, 1 month ago

Selected Answer: A

https://aws.amazon.com/blogs/big-data/query-and-visualize-aws-cost-and-usage-data-using-amazon-athena-and-amazon-quicksight/

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: A

A. Create an AWS Cost and Usage Report for the organization.

upvoted 1 times

zhangyu20000 2 years, 5 months ago

A is correct

B: no such template for cost exporer

CD: Price List Query API is for list price, not for usage

upvoted 2 times

Question #80 *Topic 1*

A company runs an IoT platform on AWS. IoT sensors in various locations send data to the company's Node.js API servers on Amazon EC2 instances running behind an Application Load Balancer. The data is stored in an Amazon RDS MySQL DB instance that uses a 4 TB General Purpose SSD volume.

The number of sensors the company has deployed in the field has increased over time, and is expected to grow significantly. The API servers are consistently overloaded and RDS metrics show high write latency.

Which of the following steps together will resolve the issues permanently and enable growth as new sensors are provisioned, while keeping this platform cost-efficient? (Choose two.)

A. Resize the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS.

B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas.

C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data.

D. Use AWS X-Ray to analyze and debug application issues and add more API servers to match the load.

E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance.

Suggested Answer: *CE*

Community vote distribution

CE (69%) BC (18%) 7%

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: CE

C and E are the correct answers.

Option C: Leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data would help to resolve the issues with the API servers being consistently overloaded. By using Kinesis, the data can be ingested and processed in real-time, allowing the API servers to handle the increased load. Using Lambda to process the data can also help to improve the overall performance and scalability of the platform.

Option E: Re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance would help to resolve the issues with high write latency. DynamoDB is a NoSQL database that is designed for high performance and scalability, making it a good fit for this use case. Additionally, DynamoDB supports auto-scaling, which can help to ensure that the database can handle the expected growth in the number of sensors.

upvoted 20 times

SuperP43 2 years, 4 months ago

I disagree with option E. Re-architecting the database tier from RDS to DynamoDB is not possible. RDS is a SQL database, and DynamoDB is a NoSQL database.

The correct one should be C and B

upvoted 9 times

ajeeshb 1 year, 3 months ago

That is why it says to "Re-architect the DB tier".

upvoted 5 times

tromyunpak 2 years ago

if it was read operations yes but the issue is write latency. also rds proxy is used to handle the write operations

upvoted 2 times

tromyunpak 2 years ago

also rds proxy is not used (sorry typo) to handle write operations properly

upvoted 1 times

kamaro 2 years, 3 months ago

I agree with you.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html
Aurora can deliver up to five times the throughput of MySQL and up to three times the throughput of PostgreSQL without requiring changes to most of your existing applications.

Aurora includes a high-performance storage subsystem. Its MySQL- and PostgreSQL-compatible database engines are customized to take advantage of that fast distributed storage. The underlying storage grows automatically as needed. An Aurora cluster volume can grow to a maximum size of 128 tebibytes (TiB).
upvoted 2 times

zejou1 2 years, 3 months ago
Naw, you can migrate: https://aws.amazon.com/blogs/big-data/near-zero-downtime-migration-from-mysql-to-dynamodb/

Plus, with DynamoDB it scales, don't need to add read replica complexity and it also supports IoT out of the box - https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SQLtoNoSQL.WhyDynamoDB.html
This is for IoT sensors that send data and I don't need to store forever so, DynamoDB for this use case is better and cheaper allowing scale
upvoted 2 times

Sarutobi 2 years, 1 month ago
I think this is the big point in this question and that DynamoDB is being position by AWS for IoT very hard. Although is technically possible to migrate with DMS from SQL to DynamoDB, is hard, but harder yet is the change of model inside the application or service.
upvoted 1 times

OCHT 2 years, 2 months ago
While options C and E may also provide some benefits, they may not address the underlying issues with the overloaded API servers and high write latency in the database. Therefore, options B and D are the best combination for resolving the issues and enabling growth as new sensors are provisioned.
upvoted 1 times

masetromain 2 years, 5 months ago
Option A, Resizing the MySQL General Purpose SSD storage to 6 TB to improve the volume's IOPS will not solve the problem, as the problem is not just related to storage size but also high write latency.

Option B, Re-architecting the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and adding read replicas would help to improve the read performance, but it won't help in reducing write latency.

Option D, Using AWS X-Ray to analyze and debug application issues and adding more API servers to match the load, would help in identifying the problem and resolving it, but it will not help in reducing the load on the servers.
upvoted 3 times

Kaps443 Most Recent 3 weeks ago
Selected Answer: CE
C + E provide the most scalable, cost-efficient, and future-proof architecture for the company's IoT platform.
upvoted 1 times

eesa 1 month, 2 weeks ago
Selected Answer: BD
B. Rediseñar a Amazon Aurora + réplicas de lectura
➔ ✅ Sí.

Aurora es mucho más escalable y eficiente que MySQL RDS normal.

Puedes crear réplicas de lectura automáticas para distribuir carga sin mucho esfuerzo manual.

Permite crecimiento masivo de forma rentable porque Aurora gestiona réplicas y escalado de manera serverless si quieres (Aurora Serverless v2).

C. Kinesis Data Streams + Lambda para ingesta
➔ ✅ Sí.

Si los sensores envían muchísimos datos en tiempo real, meterlos directo en EC2+RDS satura todo.

Kinesis puede recibir millones de eventos por segundo de forma masiva, almacenar temporalmente y procesarlos por lotes (batching) con Lambda,

desacoplando la presión sobre tus APIs y la base de datos.

Escala automáticamente y es muy rentable.

upvoted 1 times

Paul123456789 2 months, 4 weeks ago

Selected Answer: CE

A. will not fix the problem

B. read replicas will not fix the high write latency

D. is for debugging, not a solution

This make it C and E

upvoted 1 times

hhiguita 3 months, 1 week ago

Selected Answer: BC

Write performance will be improved by switch RDS to Aurora. RDS to Aurora is smooth transition without too much on the application side.

Answer E will application side not just backend DB.

upvoted 1 times

29fb203 3 months, 2 weeks ago

Selected Answer: BC

B. Re-architect the database tier to use Amazon Aurora and add read replicas

Aurora automatically scales storage up to 128 TB without manual resizing.

Faster writes and lower read latency than standard RDS MySQL.

C. Use Amazon Kinesis Data Streams and AWS Lambda for ingestion and processing

Decouples IoT data ingestion from database writes

Kinesis Data Streams ingests large volumes of sensor data without overloading API servers.

Scales automatically with the number of sensors.

Not E becvause DynamoDB is NOSql and doesn't support MySQL.

upvoted 1 times

bhanus 6 months ago

Selected Answer: BC

B. Re-architect the database tier to use Amazon Aurora instead of an RDS MySQL DB instance and add read replicas.

Amazon Aurora is a managed database service compatible with MySQL, designed for high performance and scalability.

Aurora provides better write performance and supports read replicas to handle increased read traffic as the platform grows. This will address the high write latency issue and enable horizontal scaling.

C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data.

Using Amazon Kinesis Data Streams for data ingestion offloads traffic from the API servers, reducing their load and improving scalability.

AWS Lambda can process the raw data in real time and pass it to the database or other systems, providing a cost-effective and scalable solution for data processing.

upvoted 1 times

Heman31in 6 months, 3 weeks ago

Selected Answer: CE

By combining C (Kinesis + Lambda) with E (DynamoDB), you're preparing the platform to handle exponential growth in sensor data while ensuring high availability, scalability, and low latency for both data processing and storage. This solution directly addresses the need for a robust, future-proof architecture capable of supporting massive data volumes without bottlenecks, making it well-suited for the IoT platform's growth.

upvoted 1 times

wem 6 months, 3 weeks ago

Selected Answer: BC

E would require a shift from relational to a no-sql table - what if there are multiple tables?

upvoted 1 times

konieczny69 7 months ago

Selected Answer: CE

C is straightforward.

I go for E rather than B, because db shows heavy write latency, not limit. Replacing with Aurora will speed up thing up until a limit.
Goal is to deal with it once and for all
upvoted 2 times

0b43291 7 months, 1 week ago

Selected Answer: BC

By combining options B and C, the company can address the current performance and scalability issues while enabling future growth as more sensors are deployed. Amazon Aurora provides a scalable and high-performance relational database, while Kinesis Data Streams and Lambda offer a serverless and cost-effective solution for ingesting and processing the raw data streams.

Option A may provide temporary relief by increasing IOPS, but it doesn't address the scalability and performance limitations of RDS MySQL.
Option D can help identify application issues but doesn't solve the underlying database problems.
Option E is not ideal as DynamoDB is a NoSQL database, and the existing application is likely designed for a relational database like MySQL or Aurora, requiring significant changes to the application code and data modeling.
upvoted 2 times

amministrazione 10 months ago

C. Leverage Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data.
E. Re-architect the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance.
upvoted 1 times

zolthar_z 11 months, 2 weeks ago

Selected Answer: CE

What discards B is "Add read replicas", the problem is writing the new data in the DB, adding Read replicas will increase the cost and this is not what question requests "maintain cost"
upvoted 2 times

Helpnosense 1 year ago

Selected Answer: BC

Write performance will be improved by switch RDS to Aurora. RDS to Aurora is smooth transition without too much on the application side.
Answer E will application side not just backend DB.
upvoted 2 times

TonytheTiger 1 year, 2 months ago

Selected Answer: CE

Option CE and BC. The only reason I choose E over B because said SO. Per AWS, DynamoDB is suitable for IoT (Sensor data and log ingestion)

https://docs.aws.amazon.com/whitepapers/latest/best-practices-for-migrating-from-rdbms-to-dynamodb/suitable-workloads.html
upvoted 3 times

gofavad926 1 year, 3 months ago

Selected Answer: CE

CE, kinesis + lambda & Dynamodb
upvoted 1 times

a54b16f 1 year, 3 months ago

Selected Answer: BC

Switching from RDS mysql to aurora will improve performance, by up to 10 times, which could solve the write issue. Switching from relationship database to nosql is not practical, need re-engineering whole application. plus, the performance improvement of nosql are around data read, not data write (creating/updating indexes is a huge effort)
upvoted 2 times

Question #81 *Topic 1*

A company is building an electronic document management system in which users upload their documents. The application stack is entirely serverless and runs on AWS in the eu-central-1 Region. The system includes a web application that uses an Amazon CloudFront distribution for delivery with Amazon S3 as the origin. The web application communicates with Amazon API Gateway Regional endpoints. The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket.
The company is growing steadily and has completed a proof of concept with its largest customer. The company must improve latency outside of Europe.

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

A. Enable S3 Transfer Acceleration on the S3 bucket. Ensure that the web application uses the Transfer Acceleration signed URLs.

B. Create an accelerator in AWS Global Accelerator. Attach the accelerator to the CloudFront distribution.

C. Change the API Gateway Regional endpoints to edge-optimized endpoints.

D. Provision the entire stack in two other locations that are spread across the world. Use global databases on the Aurora Serverless cluster.

E. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database.

Suggested Answer: *AC*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: AC

A and C are correct answers.
A. Enable S3 Transfer Acceleration on the S3 bucket and ensure that the web application uses the Transfer Acceleration signed URLs will accelerate the uploads of documents to S3 bucket, this will help to reduce the latency for users outside of Europe.
C. Change the API Gateway Regional endpoints to edge-optimized endpoints will help the company to improve the latency by caching the responses of the API Gateway closer to the users.
upvoted 24 times

e4bc18e 1 year, 2 months ago

A is wrong because why would you enable transfer acceleration when transfer acceleration uses the cloudfront distribution system. it makes no sense
upvoted 1 times

zolthar_z 11 months, 2 weeks ago

S3 Global acceleration is used to upload files, so, the users can upload faster the documents in any part of the world
upvoted 1 times

bcx 2 years ago

A is wrong because the users of S3 are the lambda functions, not the end user. "The API Gateway APIs call AWS Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket."
upvoted 3 times

Sab 1 year, 9 months ago

Users of S3 are not lambda, lambda is used only for writing to serverless database. Also, Aurora serverless global database only writes in one cluster and the other region cluster are used only for reads. So no matter from which location you upload, the metadata will be written to cluster in Central Europe . If it was Global DynnamoDB table then it could have helped to reduce latency.
upvoted 2 times

ninomfr64 1 year, 5 months ago

"web app uses CloudFront distribution for delivery with Amazon S3 as the origin" and "Lambda functions that store metadata in an Amazon Aurora Serverless database and put the documents into an S3 bucket" these 2 sentences let me think that users are not uploading via CluodFront into the S3 bucket at its origin, rather docs are uploaded from the Lambda
upvoted 2 times

masetromain 2 years, 5 months ago

B. Creating an accelerator in AWS Global Accelerator and attaching it to the CloudFront distribution will not help in this scenario as it only helps to route the traffic to the optimal endpoint based on the location of the user.
D. Provisioning the entire stack in two other locations that are spread across the world and using global databases on the Aurora Serverless cluster will help to reduce the latency but it would be more complex to implement and manage.
E. Adding an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database will not help in this scenario because it is only used to improve connection management and load balancing for Amazon RDS databases, but not for Aurora Serverless databases.
upvoted 5 times

masetromain 2 years, 5 months ago
https://www.examtopics.com/discussions/amazon/view/69470-exam-aws-certified-solutions-architect-professional-topic-1/
upvoted 2 times

Japanese1 1 year, 7 months ago
Complexity is not evidence against option D.
Furthermore, option D is correct because the question statement also suggests that costs can be incurred.
On the other hand, A is not a method to eliminate geographical factors.
upvoted 1 times

hussainbaloch1002 5 months, 3 weeks ago
D does not mention how to route traffic
upvoted 1 times

caputmundi666 Most Recent 2 months, 4 weeks ago
Selected Answer: CD
CD - for me. I don't understand why S3 Transfer Acceleration is better than D sincethe transfer from lambda is already on AWS's backbone.
upvoted 1 times

ParamD 3 months, 2 weeks ago
Selected Answer: AC
A is correct because it is with signed url option, lambda will facilitate signed url generation and file will be uploaded directly to S3 with transfer acceleration
upvoted 1 times

zolthar_z 11 months, 2 weeks ago
Selected Answer: AC
AC will improve latency using AWS edge locations worldwide, adding 2 locations will only benefit those 2 locations
upvoted 1 times

gfhbox0083 11 months, 4 weeks ago
A, C for sure.
B is wrong; AWS Global Accelerator and Amazon CloudFront are separate services that use the AWS global network and its edge locations around the world. CloudFront improves performance for both cacheable content (such as images and videos) and dynamic content (such as API acceleration and dynamic site delivery). Global Accelerator improves performance for a wide range of applications over TCP or UDP by proxying packets at the edge to applications running in one or more AWS Regions. Global Accelerator is a good fit for non-HTTP use cases, such as gaming (UDP), IoT (MQTT), or Voice over IP, as well as for HTTP use cases that specifically require static IP addresses or deterministic, fast regional failover. Both services integrate with AWS Shield for DDoS protection.
upvoted 2 times

red_panda 1 year, 2 months ago
Selected Answer: AC
A and C for me are the correct answers.
D is not so usefull as we are recreating the entire stack and increase a lot the costs. As first approach, A and C are the most appropriate
upvoted 2 times

failexamonly 1 year, 3 months ago
Selected Answer: AC
Aurora serverless does not support global database. search DB instance class requirements in
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database-getting-started.html
upvoted 3 times

bacharbhouri 1 year, 1 month ago
it does in V2.

[] https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v2.html#aurora-serverless-v2.advantages : Using Aurora Serverless v2 - Advantages of Aurora Serverless v2

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: AC

By elimination: B is pointless, as CF already does geo proximity. D is impossible as global DBs aren't supported by Aurora Serverless. E doesn't really help.

Remaining: A and C, which are sensible and will do the trick.

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: AC

AC, s3 transfer acceleration + edge-optimised api gateway

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: CD

This is tricky. Here is my take having in mind that the question is "The company must improve latency outside of Europe",.

A = Transfer Acceleration improves upload/downlad time, but we have already CloudFront that can also be used to speedup upload. This will not further impr
B = This will not help and also I don't know how to combine GA with CF
C = correct
D = correct
E = RDS Proxy do not improve latency

upvoted 2 times

djeong95 1 year, 4 months ago

Looks like D is wrong because you don't use global databases on the Aurora Serverless cluster. That is just not a feature given by Aurora Serverless (even global databases. "The secondary clusters" in the link below is a reference to Aurora Global Database.

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v2.how-it-works.html#aurora-serverless.ha:~:text=You%20can%20use%20Aurora%20Serverless%20v2%20capacity%20in%20the%20secondary%20clusters%20so%20they%27re%20read

upvoted 1 times

grumpysloth 6 months, 2 weeks ago

"Aurora Serverless v2 supports all manner of database workloads. Examples include development and test environments, websites, and applications tl to the most demanding, business critical applications that require high scale and high availability. It supports the full breadth of Aurora features, includ replicas. Aurora Serverless v2 is available for the Amazon Aurora MySQL-Compatible Edition and PostgreSQL-Compatible Edition."

upvoted 1 times

djeong95 1 year, 3 months ago

In addition, we are more likely to get latency from Lambda functions loading documents into S3 from API Gateway calls than we are from Lambda func

https://aws.amazon.com/blogs/compute/uploading-large-objects-to-amazon-s3-using-multipart-upload-and-transfer-acceleration/

upvoted 1 times

JMAN1 1 year, 6 months ago

Selected Answer: CD

Tricky Tricky.

A. Enable S3 Transfer Acceleration on the S3 bucket. Ensure that the web application uses the Transfer Acceleration signed URLs. -> Wrong. No such thing like TA signed URLs.
B. Create an accelerator in AWS Global Accelerator. Attach the accelerator to the CloudFront distribution. -> Wrong GA does not support CF.
C. Change the API Gateway Regional endpoints to edge-optimized endpoints.
D. Provision the entire stack in two other locations that are spread across the world. Use global databases on the Aurora Serverless cluster.
E. Add an Amazon RDS proxy between the Lambda functions and the Aurora Serverless database. -> Wrong. It is not related with latency.

upvoted 3 times

jpa8300 1 year, 5 months ago

Yes there is, https://stackoverflow.com/questions/37437782/aws-transfer-acceleration-with-pre-signed-urls-using-javascript-sdk

upvoted 1 times

JMAN1 1 year, 5 months ago

Sorry. I was wrong. Answer is A C.

serverless does not support global database and RDS proxy.

upvoted 1 times

JMAN1 1 year, 5 months ago

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.html#aurora-serverless.limitations

upvoted 1 times

bacharbhouri 1 year, 1 month ago

it does in V2.

[] https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless-v2.html#aurora-serverless-v2.advantages : Using Aurora Serverless v2 - Advantages of Aurora Serverless v2

upvoted 1 times

[Removed] 1 year, 6 months ago

Selected Answer: AC

A and C makes sense.

A is clear as what masetromain has explained.

C, An edge-optimized API endpoint typically routes requests to the nearest CloudFront Point of Presence (POP). It certaintly improve the latency of traffic originating from Europe as the traffic will now be directed to the nearest POP instead of the origin API Gateway.

upvoted 2 times

severlight 1 year, 7 months ago

Selected Answer: AC

see Sab answer

upvoted 1 times

wookchan 1 year, 8 months ago

"The company must improve latency outside of Europe."

Then in where are you going to provision an additional stack? It only says "outside of Europe."

USA? Asia? Where?

You have to consider an overall latency.

I'll go for AC

upvoted 1 times

AMohanty 1 year, 9 months ago

AD

Issue is minimize latency for "users uploading documents"

Its NOT an issue with the latency of website being delivered to the users.

Global Accelerator - Is used to decrease latency in having the user request delivered using AWS backbone network to the point of Origin

But it doesnt accelerate delivery of uploaded files into S3 .... so A is a better option.

RDS Proxy is used to decrease the time in establishing the DB connectivity ... It keeps few DB connections on warm-by condition. Option D doesn't help in reducing cross-Region latency

API Gateway edge point will reduce the latency in serving the website closer to ur location. But here question is about uploading document.

Aurora Serverless Global - can be used for uploading meta-data reducing latency time.

upvoted 1 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: AC

On a global scale, and particularly for users outside of Europe, the API Gateway and S3 access operations are the most likely components to introduce significant latency.

For the API Gateway, changing from regional endpoints to edge-optimized endpoints would bring API calls closer to global users.

For S3, enabling Transfer Acceleration would speed up the uploading and downloading of files.

Therefore, based on the provided system overview, these two components are the most likely areas needing optimization to reduce latency.

upvoted 2 times

Gabehcoud 1 year, 10 months ago

Selected Answer: CD

even though option D is complex, it would decrease the latency outside eu region.

upvoted 2 times

Question #82 *Topic 1*

An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and rafting photos and videos anytime. The photos and videos are stored in Amazon S3 Standard storage in an S3 bucket and are served through Amazon CloudFront.

The company needs to optimize the cost of the storage. A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days. The solutions architect needs to implement a solution that maintains millisecond retrieval availability of the photos and videos at the lowest possible cost.

Which solution will meet these requirements?

A. Configure S3 Intelligent-Tiering on the S3 bucket.

B. Configure an S3 Lifecycle policy to transition image objects and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days.

C. Replace Amazon S3 with an Amazon Elastic File System (Amazon EFS) file system that is mounted on Amazon EC2 instances.

D. Add a Cache-Control: max-age header to the S3 image objects and S3 video objects. Set the header to 30 days.

Suggested Answer: *B*

Community vote distribution

A (97%)

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: A

The correct answer is A. Configure S3 Intelligent-Tiering on the S3 bucket.

Amazon S3 Intelligent-Tiering is a storage class that automatically moves objects between two access tiers based on changing access patterns. Objects that are accessed frequently are stored in the frequent access tier and objects that are accessed infrequently are stored in the infrequent access tier. This allows for cost optimization without requiring manual intervention. This makes it an ideal solution for the scenario described, as it can automatically move objects that are infrequently accessed after 30 days to a lower-cost storage tier while still maintaining millisecond retrieval availability.

upvoted 16 times

masetromain 1 year, 11 months ago

Option B is not correct as it only moves data to S3 Glacier Deep Archive after 30 days, which would still require additional steps to retrieve the data.

Option C is not correct because Amazon Elastic File System (Amazon EFS) is a file storage service for use with Amazon EC2 instances, it does not provide a cost-effective solution for storing and retrieving large amounts of data.

Option D is not correct because adding a Cache-Control: max-age header only controls the caching behavior of the objects and does not address the cost optimization requirements.

upvoted 3 times

jhonivy 1 year, 11 months ago

Option D works for the reduction cost on retrieval request

upvoted 1 times

youngprinceton 1 year, 11 months ago

take the test then tell us if your answers are valid, if they are share them with us ;)

upvoted 1 times

Vsos_in29 Most Recent 10 months, 1 week ago

A is right

B S3 Glacier Deep Archive after 30 days is not correct, retrieval takes time so incorrect.

upvoted 1 times

ParamD 3 months, 2 weeks ago

Another option with S3 Glacier instant retrieval would have made the question very interesting.

upvoted 1 times

Sandeep_B 1 year, 2 months ago

Selected Answer: A

millisecond retrieval availability

upvoted 1 times

wookchan 1 year, 2 months ago

A, no brainer

upvoted 2 times

uC6rW1aB 1 year, 3 months ago

Selected Answer: A

A. Configure S3 Intelligent-Tiering on the S3 bucket: This option would automatically move objects to different storage tiers based on their access patterns. For objects that are infrequently accessed, this would help to reduce storage costs. For those that continue to be accessed frequently, they would remain in a higher-cost but faster-access tier. This should be the option that meets the requirements.

B. Configure an S3 Lifecycle policy to transition image and video objects from S3 Standard to S3 Glacier Deep Archive after 30 days: This option would significantly lower storage costs, but the retrieval time for Glacier Deep Archive could take several hours, which does not meet the millisecond retrieval requirement.

upvoted 1 times

CuteRunRun 1 year, 4 months ago

Selected Answer: A

A is right

upvoted 1 times

aviathor 1 year, 5 months ago

Selected Answer: A

B is wrong due to the Glacier Deep Archive part which is not warranted by the question.

C is wrong due to the cost of EFS and because it would require some kind of EC2 instance.

D would help caching the objects on proxies and clients, but other than that...

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

A of course

upvoted 1 times

Maria2023 1 year, 6 months ago

Selected Answer: A

I was hesitating between A and D and D looks like a really good option but it's missing one part - we do not do anything with the storage class in this option - we only update the cache TTL which would possibly reduce some costs, however, we keep paying the same price for storage. Hence I switched to A

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: A

A - easy question

upvoted 1 times

dev112233xx 1 year, 9 months ago

Selected Answer: A

A - S3 Intelligent-Tiering can fit the requirement

upvoted 1 times

God_Is_Love 1 year, 9 months ago

Selected Answer: A

First half of question drags you to answer B but SA found that some media is being used even after downloads. so data is being accessed in unknown patterns. Way to go is Intelligent tier.

upvoted 4 times

God_Is_Love 1 year, 9 months ago

*I meant even after 30 days (not downloads in above comment)

upvoted 1 times

JungMun 1 year, 10 months ago

Selected Answer: D

This is my open. The question ask us maintains millisecond retrieval ability. It means we can't use cold storage (So, A, B is not answer). EFS is expensive and not durable. If we use client cache (Ignore client's volume), we can reduce network costs(actually s3's storage costs is really cheap). It means that we can reduce costs too.

upvoted 1 times

JungMun 1 year, 10 months ago

There are lots of wrong types. Please forgive me. English is not familiar with me yet.

upvoted 2 times

c73bf38 1 year, 10 months ago

The keyword is millisecond retrieval time, which rules everything out except A.

upvoted 2 times

klog 1 year, 10 months ago

Selected Answer: A

bc A solutions architect discovers that most of the uploaded photos and videos are accessed infrequently after 30 days. However, some of the uploaded photos and videos are accessed frequently after 30 days.

upvoted 1 times

zozza2023 1 year, 11 months ago

Selected Answer: A

typico A S3 Intelligent-Tiering

upvoted 2 times

jhonivy 1 year, 11 months ago

D it will reduce the cost on retrieval requests

upvoted 1 times

Question #83 *Topic 1*

A company uses Amazon S3 to store files and images in a variety of storage classes. The company's S3 costs have increased substantially during the past year.

A solutions architect needs to review data trends for the past 12 months and identity the appropriate storage class for the objects.

Which solution will meet these requirements?

A. Download AWS Cost and Usage Reports for the last 12 months of S3 usage. Review AWS Trusted Advisor recommendations for cost savings.

B. Use S3 storage class analysis. Import data trends into an Amazon QuickSight dashboard to analyze storage trends.

C. Use Amazon S3 Storage Lens. Upgrade the default dashboard to include advanced metrics for storage trends.

D. Use Access Analyzer for S3. Download the Access Analyzer for S3 report for the last 12 months. Import the .csv file to an Amazon QuickSight dashboard.

Suggested Answer: *B*

Community vote distribution

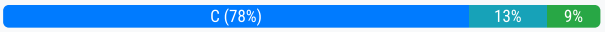

zejou1 Highly Voted 2 years, 3 months ago

Selected Answer: C

Storage class: After you configure a filter, you'll start seeing data analysis based on the filter in the Amazon S3 console in 24 to 48 hours. However, storage class analysis observes the access patterns of a filtered data set for 30 days or longer to gather information for analysis before giving a result

Storage Lens: All S3 Storage Lens metrics are retained for a period of 15 months. However, metrics are only available for queries for a specific duration, which depends on your metrics selection. This duration can't be modified. Free metrics are available for queries for a 14-day period, and advanced metrics are available for queries for a 15-month period.

You have to upgrade regardless to query up to 12 months

upvoted 15 times

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: C

Both B and C are good.

I guess AWS wants clients to use S3 Storage Lens... Hence I vote C.

upvoted 7 times

zozza2023 2 years, 5 months ago

agree with u gess aws want us to know about Lens

upvoted 3 times

gfhbox0083 Most Recent 11 months, 3 weeks ago

C, for sure.

upvoted 1 times

naylinu 1 year, 1 month ago

B ..S3 Storage Class Analysis is specifically designed to help you analyze storage access patterns. It monitors the access patterns of objects and provides insights into when it is appropriate to transition objects to different storage classes .

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, S3 Storage Lens offers comprehensive visibility into storage usage and activity trends across the AWS Organization, facilitating informed decisions on cost optimization and storage efficiency

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: C

Option C refers to using Amazon S3 Storage Lens, which provides organization-wide visibility into object storage usage and activity trends. By upgrading to include advanced metrics and recommendations, users can access detailed insights that help optimize storage costs across their S3 resources. S3 Storage Lens offers dashboard views and metrics that can directly inform on the appropriate storage class based on actual usage patterns, making it a comprehensive solution for the stated requirements.

upvoted 1 times

AWSPro1234 1 year, 5 months ago

Amazon S3 Storage Class Analysis:

Amazon S3 provides a Storage Class Analysis tool that helps you analyze access patterns to your S3 objects over time. You can enable it on your S3 bucket to collect data on object access patterns.

upvoted 1 times

AWSPro1234 1 year, 5 months ago

Answer is B.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

To me here the key sentence is "review data trends for the past 12 months"

A = CUR provides detailed usage data but it is not the best tool for this job

B = S3 storage class analysis provides recommendation for Standard and Standard IA storage classes, but does not provide data trends

C = correct

D = Access Analyzer provides visibility for buckets that are configured to allow access to anyone on the internet or other AWS accounts

upvoted 1 times

Nicoben 1 year, 6 months ago

Selected Answer: B

B is the right answer, because it suffices a bucket analysis --> https://docs.aws.amazon.com/AmazonS3/latest/userguide/analytics-storage-class.html

C instead is a solution for a more organization-wide analysis of bucket --> https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens.html

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: C

see AMohanty answer

upvoted 1 times

Simon523 1 year, 9 months ago

Selected Answer: B

S3 Storage Class Analysis enables you to monitor access patterns across objects to help you decide when to transition data to the right storage class to optimize costs.

upvoted 2 times

jpa8300 1 year, 5 months ago

Storage Class is only used for recommendation for Standard to Standard IA

upvoted 1 times

AMohanty 1 year, 9 months ago

C

Storage Class is only used for recommendation for Standard to Standard IA

upvoted 4 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: C

Option B: Amazon S3's Storage Class Analysis function is mainly used to analyze the access patterns of objects in S3 buckets so that you can transfer these objects to the most cost-effective storage class. However, this feature does not provide detailed historical data for the past 12 months; it is more about observing current usage patterns and making the best storage class decisions based on those patterns.

If you need detailed storage trends and object status over the past 12 months, option C (using Amazon S3 Storage Lens) may be a better choice.

Amazon S3 Storage Lens provides comprehensive storage analysis, including historical trends and advanced metrics, which may be more suitable for analyzing long-term data and storage conditions.

upvoted 2 times

YodaMaster 1 year, 12 months ago

I choose C.

B. Storage class analysis only provides recommendations for Standard to Standard IA classes. The company uses a variety of storage classes.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

a hard one ... I guess C, but could be B :/

upvoted 1 times

Limlimwdwd 2 years ago

Selected Answer: B

By using Amazon S3 analytics Storage Class Analysis you can analyze storage access patterns to help you decide when to transition the right data to the right storage class. This new Amazon S3 analytics feature observes data access patterns to help you determine when to transition less frequently accessed STANDARD storage to the STANDARD_IA (IA, for infrequent access) storage class.

So it meet the qn objective of identify the appropriate storage class for the objects

upvoted 1 times

Question #84 *Topic 1*

A company has its cloud infrastructure on AWS. A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts.

What should the solutions architect do to meet these requirements?

A. Use AWS CloudFormation templates. Add IAM policies to control the various accounts, Deploy the templates across the multiple Regions.

B. Use AWS Organizations. Deploy AWS CloudFormation templates from the management account Use AWS Control Tower to manage deployments across accounts.

C. Use AWS Organizations and AWS CloudFormation StackSets. Deploy a Cloud Formation template from an account that has the necessary IAM permissions.

D. Use nested stacks with AWS CloudFormation templates. Change the Region by using nested stacks.

Suggested Answer: *C*

Community vote distribution

C (100%)

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: C

The correct answer is C. Use AWS Organizations and AWS CloudFormation StackSets.
AWS Organizations allows the management of multiple AWS accounts as a single entity and AWS CloudFormation StackSets allows creating, updating, and deleting stacks across multiple accounts and regions in an organization. This solution allows creating a single CloudFormation template that can be deployed across multiple accounts and regions, and also allows for the management of access and permissions for the different accounts through the use of IAM roles and policies in the management account.
upvoted 17 times

masetromain 1 year, 11 months ago

Option A and D both use AWS CloudFormation, but do not take into account the management of multiple accounts and regions. Option B uses AWS Organizations but doesn't include the use of CloudFormation StackSets, which is necessary for managing deployments across multiple accounts and regions.
upvoted 6 times

jpa8300 12 months ago

I agree with what you say here, C is a good choice, but in B they mention Control Tower which is also used to manage multiple accounts, couldn't it be a correct answer also?
upvoted 1 times

ninomfr64 Most Recent 11 months, 2 weeks ago

Selected Answer: C

A = cloud work but it is hard
B = Control Tower cannot manage stack deployments across accounts
C = correct
D = nested stack allows to provision resources by using different CloudFormation templates
upvoted 3 times

totten 1 year, 2 months ago

Selected Answer: C

Option C is the most suitable. Here's why:

AWS Organizations: AWS Organizations helps you centrally manage multiple AWS accounts, which is especially useful when dealing with multiple Regions and accounts. You can organize your accounts into an organizational structure, apply policies across accounts, and manage billing.

AWS CloudFormation StackSets: StackSets is a CloudFormation feature that enables you to deploy CloudFormation stacks across multiple accounts and Regions with a single CloudFormation template. This simplifies the process of deploying and managing infrastructure consistently across your organization.
upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

C no doubt

upvoted 2 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: C

keywords = AWS Organizations && AWS CloudFormation StackSets.

upvoted 1 times

rbm2023 1 year, 7 months ago

Selected Answer: C

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

Cloud Formation Stack Sets allow you to roll out Cloud Formation stacks over multiple AWS accounts and in multiple Regions with just a couple of clicks. When we launched Stack Sets, grouping accounts was primarily for billing purposes. Since the launch of AWS Organizations, you can centrally manage multiple AWS accounts across diverse business needs including billing, access control, compliance, security, and resource sharing.

upvoted 3 times

mfsec 1 year, 9 months ago

Selected Answer: C

Use AWS Organizations and AWS CloudFormation StackSets

upvoted 2 times

zozza2023 1 year, 11 months ago

Selected Answer: C

The correct answer is C

upvoted 4 times

Question #85 *Topic 1*

A company has its cloud infrastructure on AWS. A solutions architect needs to define the infrastructure as code. The infrastructure is currently deployed in one AWS Region. The company's business expansion plan includes deployments in multiple Regions across multiple AWS accounts.

What should the solutions architect do to meet these requirements?

A. Use AWS CloudFormation templates. Add IAM policies to control the various accounts, Deploy the templates across the multiple Regions.

B. Use AWS Organizations. Deploy AWS CloudFormation templates from the management account Use AWS Control Tower to manage deployments across accounts.

C. Use AWS Organizations and AWS CloudFormation StackSets. Deploy a Cloud Formation template from an account that has the necessary IAM permissions.

D. Use nested stacks with AWS CloudFormation templates. Change the Region by using nested stacks.

Suggested Answer: *C*

Community vote distribution

C (100%)

masetromain Highly Voted 1 year, 5 months ago
same question of "Questions #84"
upvoted 17 times

yorkicurke Highly Voted 8 months ago
These Site Moderators getting lazy boy!
upvoted 6 times

NikkyDicky Most Recent 12 months ago
Selected Answer: C
C. a dup question
upvoted 2 times

rbm2023 1 year, 1 month ago
Selected Answer: C
This question is duplicated in the Exam Topics site. Question 85 is the same as Question 84
upvoted 1 times

bordy20 1 year, 1 month ago
C:
https://sanderknape.com/2017/07/cloudformation-stacksets-automated-cross-account-region-deployments/#:~:text=A%20StackSet%20is%20a%20set,deploying%20to%20multiple%20accounts%2Fregions.
upvoted 1 times

Nguyen25183 1 year, 2 months ago
Thought that my internet was intertupted. then i was wrong =)))
upvoted 4 times

Musk 1 year, 4 months ago
This is repeated :-(
upvoted 2 times

tatdatpham 1 year, 4 months ago
Selected Answer: C
Duplicate question with #84
upvoted 3 times

zhangyu20000 1 year, 5 months ago
C is correct answer
upvoted 3 times

Question #86 *Topic 1*

A company plans to refactor a monolithic application into a modern application design deployed on AWS. The CI/CD pipeline needs to be upgraded to support the modern design for the application with the following requirements:

• It should allow changes to be released several times every hour.
• It should be able to roll back the changes as quickly as possible.

Which design will meet these requirements?

A. Deploy a CI/CD pipeline that incorporates AMIs to contain the application and their configurations. Deploy the application by replacing Amazon EC2 instances.

B. Specify AWS Elastic Beanstalk to stage in a secondary environment as the deployment target for the CI/CD pipeline of the application. To deploy, swap the staging and production environment URLs.

C. Use AWS Systems Manager to re-provision the infrastructure for each deployment. Update the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and use Amazon Route 53 weighted routing to point to the new environment.

D. Roll out the application updates as part of an Auto Scaling event using prebuilt AMIs. Use new versions of the AMIs to add instances. and phase out all instances that use the previous AMI version with the configured termination policy during a deployment event.

Suggested Answer: *B*

Community vote distribution

B (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct answer is B. Specifying AWS Elastic Beanstalk to stage in a secondary environment as the deployment target for the CI/CD pipeline of the application and swapping the staging and production environment URLs. This approach allows the company to deploy updates several times an hour and quickly roll back changes as needed.

Option A, Deploying a CI/CD pipeline that incorporates AMIs to contain the application and their configurations. Deploy the application by replacing Amazon EC2 instances, while it may provide a way to roll back changes by replacing instances with previous versions, it may not allow for rapid deployment of updates multiple times per hour.

upvoted 18 times

masetromain 2 years, 5 months ago

Option C, Using AWS Systems Manager to re-provision the infrastructure for each deployment. Updating the Amazon EC2 user data to pull the latest code artifact from Amazon S3 and using Amazon Route 53 weighted routing to point to the new environment, would require more time-consuming steps and may not be able to roll back changes as quickly.

Option D, Rolling out the application updates as part of an Auto Scaling event using prebuilt AMIs. Using new versions of the AMIs to add instances and phasing out all instances that use the previous AMI version with the configured termination policy during a deployment event, while it may be a way to roll back changes, it doesn't allow for rapid deployment of updates multiple times per hour.

upvoted 5 times

jpa8300 1 year, 5 months ago

Good explanation, but concerning option C it is not quite right, you say that 'may not be able to roll back changes as quickly.', but since it is using Route 53 weighted configuration, in case of failure of the new instances, you just need to change again the weighted configuration to point 100% to the old instances while you replace again the new instances by old instances.

upvoted 1 times

gfhbox0083 Most Recent 11 months, 3 weeks ago

Selected Answer: B

B, for sure.
Using AWS Elastic Beanstalk environment Swap.
https://docs.aws.amazon.com/whitepapers/latest/blue-green-deployments/swap-the-environment-of-an-elastic-beanstalk-application.html

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: B

A = replacing existing EC2 instances does not allow for roll back the changes as quickly as possible

B = correct (tough Beanstalk is not the best service for releasing several times every hour)

C = could work, but here you are combining SSM and user data to achieve what beanstalk does natively

D = this would not work as you need to build AMIs (AMI Builder not mentioned) and also rapid rollback is better achieved avoiding termination of old AMI version

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

probably B

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: B

Imagine the cost for replacing AMIs and EC2 or re-provision infrastructure several times per day. Although cost effectiveness is not part the requirement in the question. the only option that seems correct is B.

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

B. Specify AWS Elastic Beanstalk

upvoted 1 times

Untamables 2 years, 5 months ago

Selected Answer: B

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html

upvoted 3 times

Question #87 *Topic 1*

A company has an application that runs on Amazon EC2 instances. A solutions architect is designing VPC infrastructure in an AWS Region where the application needs to access an Amazon Aurora DB Cluster. The EC2 instances are all associated with the same security group. The DB cluster is associated with its own security group.

The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB Cluster.

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

A. Add an inbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the source over the default Aurora port.

B. Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port.

C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port.

D. Add an outbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the destination over the default Aurora port.

E. Add an outbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the destination over the ephemeral ports.

Suggested Answer: *AB*

Community vote distribution

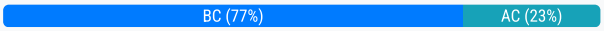

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: BC

The correct combination of steps to meet these requirements is B and C.

B. Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port. This allows the instances to make outbound connections to the DB cluster on the default Aurora port.

C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port. This allows connections to the DB cluster from the EC2 instances on the default Aurora port.

upvoted 32 times

masetromain 2 years, 5 months ago

A. Adding an inbound rule to the EC2 instances' security group would allow incoming connections to the instances on the default Aurora port, but it would not allow the instances to connect to the DB cluster.

D. Adding an outbound rule to the DB cluster's security group would allow the DB cluster to make outbound connections to the EC2 instances on the default Aurora port, but it would not allow connections to the DB cluster from the instances.

E. Adding an outbound rule to the DB cluster's security group specifying the EC2 instances' security group as the destination over the ephemeral ports would allow the DB cluster to make outbound connections to the instances on ephemeral ports, but it would not allow connections to the DB cluster from the instances on the default Aurora port.

upvoted 3 times

vjp_training 1 year, 9 months ago

Security group is stateful. So you just need to set up inbound

upvoted 3 times

HussamShokr 2 years ago

why we should add an outbound rule to the EC2 instances' security group??? it is already allowed by default in the EC2 security group becauce all outbound ports are allowed by default.

upvoted 3 times

jainparag1 1 year, 7 months ago

wow..then in that case your EC2 instance can talk to anything. No SG rule is required. You need to establish a connectivity route first.

upvoted 1 times

ninomfr64 1 year, 5 months ago

it is the other way around, all connection are denied and you can only allow connection. You need outbound from EC2 to Aurora to allow the app initiate a connection to the database instance

upvoted 2 times

c73bf38 Highly Voted 2 years, 4 months ago

Selected Answer: AC

To provide the application with least privilege access to the Aurora DB cluster, the solutions architect should add inbound rules to both the security groups.

For the EC2 instances' security group, an inbound rule should be added that allows traffic from the DB cluster's security group over the default Aurora port. This will allow the EC2 instances to communicate with the Aurora DB cluster.

For the Aurora DB cluster's security group, an inbound rule should be added that allows traffic from the EC2 instances' security group over the default Aurora port. This will allow the Aurora DB cluster to communicate with the EC2 instances.

By default all outbound rules are open, it's only the ingress that needs to allow traffic.

upvoted 12 times

c73bf38 2 years, 4 months ago

B&C after doing a recreate in the AWS Console, stand corrected.

upvoted 7 times

c73bf38 2 years, 4 months ago

To provide the application with least privilege access to the Amazon Aurora DB Cluster, the solutions architect should take the following steps:

Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port (port 3306). This will allow the EC2 instances to connect to the Aurora DB Cluster.

Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port (port 3306). This will allow the EC2 instances to send traffic to the Aurora DB Cluster.

upvoted 3 times

penguins2 Most Recent 7 months, 3 weeks ago

BC. The steps are clearly stated here: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/tutorial-ec2-rds-option3.html#option3-task3-connect-rds-database-to-ec2-instance

upvoted 1 times

nimbus_00 8 months, 3 weeks ago

Selected Answer: BC

NB. The DB cluster doesn't need to initiate connections to the EC2 instances.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: BC

BC, ec2 -> bd; ec2 outbound rule to allow access to bd; db inbound rule to allow access from ec2

upvoted 2 times

igor12ghsj577 1 year, 4 months ago

Selected Answer: BC

Tricky question. They say with least privileges, so I think they don't want to use default (allow-all) rule, but limit as much as possible and allow only specific traffic to DB)

"By default, a security group includes an outbound rule that allows all outbound traffic. We recommend that you remove this default rule and add outbound rules that allow specific outbound traffic only."

https://docs.aws.amazon.com/quicksight/latest/user/vpc-security-groups.html

upvoted 3 times

cox1960 1 year, 5 months ago

CE

- A and B are nonsense, since they talk about aurora port on ec2 SGs. In SG you always put rules on the local ports.

- C obvious
- E over D, always ephemeral on outbound, but at the condition we replace the existing all open rule
upvoted 1 times

jpa8300 1 year, 5 months ago

Selected Answer: BC

I believe that C is enough, we don't need to define the outbound from EC2 to DB, but since we have to choose two, the only other option that is correct is B. And someone say below that have tested this configuration, so I hope he tested defining only what is mentioned in C, to see if it is enough or not. It would be nice.
upvoted 3 times

shaaam80 1 year, 6 months ago

Selected Answer: BC

Answer - B& C
Outbound rule to the EC2 SG with DB SG as destination
Inbound rule to the DB SG with EC2 SG as source
upvoted 1 times

eurriola10 1 year, 7 months ago

Selected Answer: AC

Security Groups are stateful, that means you don't need to specify an outbound rule if you have an inbound rule that permit access to the resource.
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-security-groups.html#security-group-basics

In other hand, the outbound traffic rules typically don't apply to DB clusters. Outbound traffic rules apply only if the DB cluster acts as a client.
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.RDSSecurityGroups.html#Overview.RDSSecurityGroups.VPCSec.

Because of that B, D and E are wrong answers
upvoted 1 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: AC

By default, AWS Security Groups allow all outbound traffic. Therefore, in most cases, there's no need to configure outbound rules unless you have specific security requirements.

Add an inbound rule to the EC2 instance's security group, setting the DB cluster's security group as the source over Aurora's default port. This enables interaction between the DB Cluster and the EC2 instances. Corresponds to Option A.

Add an inbound rule to the DB Cluster's security group, setting the EC2 instance's security group as the source over Aurora's default port. This allows the EC2 instances to interact with the DB Cluster. Corresponds to Option C.
upvoted 2 times

- **uC6rW1aB** 1 year, 9 months ago
  By the way, the outbound rules are unnecessary in this case because the database cluster does not need to access any data from the application. The database cluster only needs to receive traffic from the application so that the application can read and write to the database.
  upvoted 1 times

- **eurriola10** 1 year, 7 months ago
  my two cents.
  Agree AC are the correct answer.
  Security Groups are stateful, that means you don't need to specify an outbound rule if you have an inbound rule that permit access to the resource.
  https://docs.aws.amazon.com/vpc/latest/userguide/vpc-security-groups.html#security-group-basics
  In other hand, the outbound traffic rules typically don't apply to DB clusters. Outbound traffic rules apply only if the DB cluster acts as a client.
  https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Overview.RDSSecurityGroups.html#Overview.RDSSecurityGroups.VPCSec
  upvoted 1 times

vjp_training 1 year, 10 months ago

Selected Answer: AC

By default, all outbound rules are allow
upvoted 1 times

- **vn_thanhtung** 1 year, 10 months ago
  Don't provide wrong answer. Answer is B,C
  upvoted 1 times

jainparag1 1 year, 7 months ago

you are providing the wrong answer. The correct answer is AC. Inbound rules are supposed to be added.

upvoted 1 times

vn_thanhtung 1 year, 10 months ago

The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB Cluster.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: BC

BC of course

upvoted 2 times

jainparag1 1 year, 7 months ago

AC is correct.

upvoted 1 times

bcx 2 years ago

Selected Answer: BC

It is outbound from the clients to the db server listening port. And inbound to the db server listening ports from the clients.

upvoted 2 times

Jonalb 2 years ago

Selected Answer: BC

"My choice relays on the fact that the security groups are stateful, so we only need to allow the outbound traffic for the ec2 instances to pass and the return will also be allowed. Same for the RDS. This combination is also based on the standard traffic flow initiated from instance to DB"

upvoted 1 times

Maria2023 2 years, 2 months ago

Selected Answer: BC

My choice relays on the fact that the security groups are stateful, so we only need to allow the outbound traffic for the ec2 instances to pass and the return will also be allowed. Same for the RDS. This combination is also based on the standard traffic flow initiated from instance to DB.

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: BC

BC gets my vote

upvoted 2 times

Question #88 *Topic 1*

A company wants to change its internal cloud billing strategy for each of its business units. Currently, the cloud governance team shares reports for overall cloud spending with the head of each business unit. The company uses AWS Organizations to manage the separate AWS accounts for each business unit. The existing tagging standard in Organizations includes the application, environment, and owner. The cloud governance team wants a centralized solution so each business unit receives monthly reports on its cloud spending. The solution should also send notifications for any cloud spending that exceeds a set threshold.

Which solution is the MOST cost-effective way to meet these requirements?

A. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in each account to create monthly reports for each business unit.

B. Configure AWS Budgets in the organization's management account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's management account to create monthly reports for each business unit.

C. Configure AWS Budgets in each account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit.

D. Enable AWS Cost and Usage Reports in the organization's management account and configure reports grouped by application, environment. and owner. Create an AWS Lambda function that processes AWS Cost and Usage Reports, sends budget alerts, and sends monthly reports to each business unit's email list.

Suggested Answer: *B*

Community vote distribution

masetromain Highly Voted 1 year, 5 months ago

Selected Answer: B

B. Configure AWS Budgets in the organization's management account and configure budget alerts that are grouped by application, environment, and owner. Add each business unit to an Amazon SNS topic for each alert. Use Cost Explorer in the organization's management account to create monthly reports for each business unit.

This option is the most cost-effective because it utilizes the organization's management account to set budgets and configure alerts for all accounts in the organization, rather than having to configure budgets and alerts individually in each account. Additionally, using Cost Explorer in the management account allows the cloud governance team to view the consolidated spending for all accounts in the organization and create reports for each business unit. This eliminates the need to access each individual account to view costs and create reports.

upvoted 25 times

masetromain 1 year, 5 months ago

Option A is not the most cost-effective solution because it requires configuring budgets and reports in multiple accounts, which increases the complexity and cost of managing the cloud spending for each business unit.

Option C is not the most cost-effective solution because it requires the cloud governance team to access the AWS Billing and Cost Management dashboard in each account to create monthly reports for each business unit, which increases the complexity and cost of managing the cloud spending for each business unit.

Option D is not the most cost-effective solution because it requires creating an AWS Lambda function to process AWS Cost and Usage Reports, which increases the complexity and cost of managing the cloud spending for each business unit.

upvoted 6 times

NikkyDicky Most Recent 12 months ago

Selected Answer: B

B for sure

upvoted 1 times

Maria2023 1 year ago

"configure budget alerts that are grouped by application, environment, and owner" - I just literally tried to create a budget alert and I am not able to see any option for grouping by tags. Another nonsense question

upvoted 2 times

b3llman 10 months, 3 weeks ago

Billing > Budgets > Create budget > Customize (advanced) > Budget scope > Filter specific AWS cost dimensions

upvoted 1 times

SkyZeroZx 1 year ago

Selected Answer: B

keyword = AWS Budgets in the organization's management

other more overhead each by account

upvoted 2 times

yama234 1 year, 2 months ago

B

centralized solution = management account

send notifications for any cloud spending that exceeds a set threshold = AWS Budgets

https://aws.amazon.com/blogs/mt/manage-cost-overruns-part-1/

upvoted 4 times

mfsec 1 year, 3 months ago

Selected Answer: B

B. Configure AWS Budgets in the organization's management account

upvoted 1 times

Question #89 *Topic 1*

A company is using AWS CloudFormation to deploy its infrastructure. The company is concerned that, if a production CloudFormation stack is deleted, important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted.

How can the company prevent users from accidentally deleting data in this way?

A. Modify the CloudFormation templates to add a DeletionPolicy attribute to RDS and EBS resources.

B. Configure a stack policy that disallows the deletion of RDS and EBS resources.

C. Modify IAM policies lo deny deleting RDS and EBS resources that are tagged with an "aws:cloudformation:stack-name" tag.

D. Use AWS Config rules to prevent deleting RDS and EBS resources.

Suggested Answer: *A*

Community vote distribution

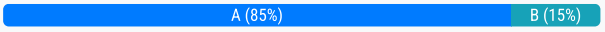

zejou1 Highly Voted 2 years, 3 months ago

Selected Answer: A

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html
With the DeletionPolicy attribute you can preserve, and in some cases, backup a resource when its stack is deleted. You specify a DeletionPolicy attribute for each resource that you want to control. If a resource has no DeletionPolicy attribute, AWS CloudFormation deletes the resource by default.
Retain
CloudFormation keeps the resource without deleting the resource or its contents when its stack is deleted. You can add this deletion policy to any resource type. When CloudFormation completes the stack deletion, the stack will be in Delete_Complete state; however, resources that are retained continue to exist and continue to incur applicable charges until you delete those resource

upvoted 17 times

nimbus_00 Most Recent 8 months, 3 weeks ago

Selected Answer: A

By adding the DeletionPolicy attribute to the CloudFormation template for RDS and EBS resources, you can specify actions to be taken when a stack is deleted. Setting the DeletionPolicy to Retain ensures that the RDS and EBS resources are not deleted when the CloudFormation stack is deleted.

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: A

Option A is the correct approach because CloudFormation allows you to specify a DeletionPolicy attribute for resources within your templates. This attribute can prevent resources like Amazon RDS databases and Amazon EBS volumes from being deleted when the stack is deleted. You can set the DeletionPolicy to “Retain” for specific resources, ensuring they are not automatically removed alongside the stack.

upvoted 1 times

Maygam 1 year, 6 months ago

Selected Answer: B

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack-resources.html

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A, basic DeletionPolicy use case

upvoted 2 times

aviathor 1 year, 10 months ago

Yes but should be supplemented with deletion protection on the database.

upvoted 2 times

Maria2023 2 years ago

Selected Answer: A

Although that I would preferably use both A and B - this is an exam and the truth is in the wording - "important data stored in Amazon RDS databases or Amazon EBS volumes might also be deleted" - we don't care if the resources are deleted but the data, which makes me believe they want us to set up a deletion policy at a resource level to "Retain"

upvoted 2 times

zak340 2 years ago

Selected Answer: B

Explanation:
Stack policies are a powerful feature of AWS CloudFormation that allows you to control fine-grained permissions for resources within a stack. By configuring a stack policy that disallows the deletion of RDS and EBS resources, you can prevent users from accidentally deleting these critical resources and the associated data.

Option A (Modifying CloudFormation templates with DeletionPolicy attribute) is not the best solution in this case. While the DeletionPolicy attribute can be used to control resource behavior during stack deletion, it is not applicable to Amazon RDS instances or Amazon EBS volumes.

upvoted 2 times

fartosh 1 year, 1 month ago

> the DeletionPolicy attribute [...] is not applicable to Amazon RDS instances or Amazon EBS volumes.
This statement is false. From https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html
Retain
[...] You can add this deletion policy to any resource type.

Snapshot
Resources that support snapshots include:
[...]
- AWS::EC2::Volume
[...]
- AWS::RDS::DBInstance

upvoted 1 times

bcx 2 years ago

The correct answer is A, not because what you say is wrong, but because the question states that the stacks can be deleted, you cannot prevent the deletion of the stack (as required by the question). So the DeletionPolicy will let you delete the stack and retain or take a snapshot of the Database/BUCKET/... (whichever is applicable). You will not lose any data in that case and the stack would have been succesfully deleted.

upvoted 3 times

rbm2023 2 years, 1 month ago

Selected Answer: A

Check the differences and use cases where to use a stack policy or add a deletion policy (retain):
Stack policy and deletion policy are both ways to protect resources created by CloudFormation stacks, but they have different functions.
Stack policy is a feature that allows you to specify a JSON policy document that restricts what actions can be taken on a CloudFormation stack. Stack policies are used to prevent accidental or intentional updates or deletions of critical resources in your stack, by specifying which resources can be modified and by whom. Stack policies can be used to allow specific teams or individuals to modify specific resources in a stack while preventing them from modifying others.

upvoted 3 times

rbm2023 2 years, 1 month ago

Deletion policy, on the other hand, is a property of certain AWS resources that determines what happens to the resource when the stack is deleted. The deletion policy can be set to one of three values: "Delete", "Retain", or "Snapshot". When the deletion policy is set to "Delete", the resource is deleted when the stack is deleted. When the deletion policy is set to "Retain", the resource is not deleted when the stack is deleted, but must be deleted manually. When the deletion policy is set to "Snapshot", the resource is deleted when the stack is deleted, but a snapshot of the resource is retained.
In summary, stack policies are used to control what changes can be made to a stack, while deletion policies are used to determine what happens to resources when a stack is deleted.

upvoted 1 times

OCHT 2 years, 2 months ago

Selected Answer: B

ption B, which suggests configuring a stack policy that disallows the deletion of RDS and EBS resources, is better in this scenario. While using DeletionPolicy attribute (Option A) can be helpful for preserving and backing up the resource, it does not address the problem of accidental deletion of resources or control access to delete the resource.

On the other hand, a Stack Policy can be used to prevent accidental deletion of resources by specifying which actions can be performed on the resources within in the stack, thereby adding an essential layer of protection.

By implementing a Stack Policy, a company can limit updating the resources in the stack, control who can make changes to the stack, and prevent accidental deletion of resources. Therefore, configuring a Stack Policy is necessary and more satisfactory to protect data from accidental deletion while using AWS CloudFormation.
upvoted 1 times

Sarutobi 2 years, 2 months ago
You are correct about the process of the UPDATE stack action. What happens to the resources created by the CloudFormation stack when the stack itself is deleted?
upvoted 1 times

mfsec 2 years, 3 months ago
Selected Answer: A
A for sure
upvoted 2 times

kiran15789 2 years, 3 months ago
Selected Answer: B
A stack policy is a document that defines the update and deletion actions that can be performed on resources in a CloudFormation stack. By default, all resources in a CloudFormation stack can be deleted by users with appropriate permissions. However, you can use a stack policy to restrict the deletion of certain resources, such as Amazon RDS databases or Amazon EBS volumes.

In this case, the company can create a stack policy that explicitly disallows the deletion of any RDS or EBS resources in the production CloudFormation stack. This will prevent users from accidentally deleting important data stored in these resources.
upvoted 1 times

God_Is_Love 2 years, 3 months ago
Selected Answer: A
For RDS instances, you can set the "DeletionPolicy" attribute to "Retain". This will ensure that when the stack is deleted, the RDS instance will not be deleted and its data will be retained.

For EBS volumes, you can use the "DeletionPolicy" attribute in combination with the "SnapshotId" attribute to create a snapshot of the volume before deleting it. This will allow you to restore the data later if need

Yaml examples for RDS and EBS :

Resources:
MyDB:
Type: AWS::RDS::DBInstance
Properties:
RDS instance properties go here
DeletionPolicy: Retain

Resources:
MyVolume:
Type: AWS::EC2::Volume
Properties:
Volume properties go here
DeletionPolicy: Snapshot
SnapshotId: my-snapshot-id
upvoted 1 times

spd 2 years, 4 months ago
Selected Answer: A
Clear A
upvoted 1 times

lunt 2 years, 4 months ago
Selected Answer: A

AC1984 do your homework.
Stack policy can protect against deletion but not against actual entire CFN stack template being deleted. DeletionPolicy = if I was to delete the entire CFN stack, the CFN process will delete all elements and skip over RDS and EBS due to protections. 20 second Google search could of confirmed this.
upvoted 2 times

AC1984 2 years, 4 months ago
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/protect-stack-resources.html
upvoted 1 times

AC1984 2 years, 4 months ago
Selected Answer: B
B. Configure a stack policy that disallows the deletion of RDS and EBS resources.
A stack policy is a JSON-based document that defines the actions that can be performed on a CloudFormation stack, and can be used to prevent users from accidentally deleting critical resources. By configuring a stack policy that disallows the deletion of RDS and EBS resources, the company can prevent users from accidentally deleting important data stored in those resources.

Option A (adding a DeletionPolicy attribute) does not prevent users from deleting the resources, but rather determines what happens to the resources when the stack is deleted. Option C (modifying IAM policies) is not sufficient because it only affects the permissions of specific users or groups, and does not prevent accidental deletions. Option D (using AWS Config rules) can help detect deletions of RDS and EBS resources, but it does not prevent them from being deleted.
upvoted 1 times

sambb 2 years, 3 months ago
"Option A (adding a DeletionPolicy attribute) does not prevent users from deleting the resources, but rather determines what happens to the resources when the stack is deleted." This is actually what the question is asking !
upvoted 1 times

moota 2 years, 4 months ago
Selected Answer: A
I go for A because I assume that the CF stack is allowed to be deleted in some deployment scenarios.
upvoted 1 times

Question #90 *Topic 1*

A company has VPC flow logs enabled for Its NAT gateway. The company is seeing Action = ACCEPT for inbound traffic that comes from public IP address 198.51.100.2 destined for a private Amazon EC2 instance.

A solutions architect must determine whether the traffic represents unsolicited inbound connections from the internet. The first two octets of the VPC CIDR block are 203.0.

Which set of steps should the solutions architect take to meet these requirements?

A. Open the AWS CloudTrail console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interlace. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

B. Open the Amazon CloudWatch console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

C. Open the AWS CloudTrail console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

D. Open the Amazon CloudWatch console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 198.51.100.2" and the source address set as "like 203.0". Run the stats command to filter the sum of bytes transferred by the source address and the destination address.

Suggested Answer: *D*

Community vote distribution

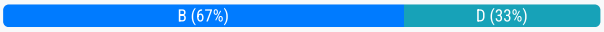

vsk12 Highly Voted 2 years, 5 months ago

I would go with option B. Source will be public IP like 198.51.100.2.

upvoted 22 times

kiran15789 Highly Voted 2 years, 3 months ago

Selected Answer: B

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/

Refer Reason 1

Run the query below.
filter (dstAddr like 'xxx.xxx' and srcAddr like 'public IP')
| stats sum(bytes) as bytesTransferred by srcAddr, dstAddr
| limit 10
Note: You can use just the first two octets in the search filter to analyze all network interfaces in the VPC. In the example above, replace xxx.xxx with the first two octets of your VPC classless inter-domain routing (CIDR). Also, replace public IP with the public IP that you're seeing in the VPC flow log entry.

Query results show traffic on the NAT gateway private IP from the public IP, but not traffic on other private IPs in the VPC. These results confirm that the incoming traffic was unsolicited. However, if you do see traffic on the private instance's IP, then follow the steps under Reason #2.

upvoted 19 times

sashenka 7 months, 3 weeks ago

To determine whether the traffic represents unsolicited inbound connections from the internet, use the Amazon CloudWatch console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". This approach helps you analyze the VPC flow logs to identify if the inbound traffic to the private EC2 instance is expected return traffic or unsolicited. The stats command can be used to filter the sum of bytes transferred by the source address and the destination address, providing insight into the traffic patterns and ensuring network security.

upvoted 1 times

sashenka 7 months, 3 weeks ago

It has to be D as it only includes unsolicited traffic. Option B includes both.

upvoted 1 times

zejou1 2 years, 3 months ago

For those that are choosing D - this is why D is incorrect and needs to be B

upvoted 2 times

papan83 Most Recent 1 month, 3 weeks ago

Selected Answer: D

Step-by-Step Approach:

You already have this log entry:

srcaddr = 198.51.100.2
dstaddr = 203.0.x.x
action = ACCEPT

Now, query CloudWatch Logs for the reverse flow:

srcaddr = 203.0.x.x (your EC2 instance)

dstaddr = 198.51.100.2 (the public IP)

If you find outbound traffic from your instance to the public IP before the inbound traffic, then the connection is solicited (i.e. it's a reply).

If you do not find any outbound flow to that public IP, then the traffic is unsolicited — a potentially unexpected or malicious inbound attempt that should have been blocked.

upvoted 1 times

BennyMao 3 months, 3 weeks ago

Selected Answer: D

The NAT gateway allows outbound internet traffic from private instances but does not accept unsolicited inbound connections.
If 198.51.100.2 is contacting the private instance, we need to determine if this is a response to an existing outbound request from the private instance.

upvoted 1 times

BennyMao 3 months, 3 weeks ago

Selected Answer: D

The NAT gateway allows outbound internet traffic from private instances but does not accept unsolicited inbound connections.
If 198.51.100.2 is contacting the private instance, we need to determine if this is a response to an existing outbound request from the private instance.

upvoted 1 times

grumpysloth 6 months, 2 weeks ago

Selected Answer: D

we need to check if the request starts from ec2 instances outbound, not the other way round.

upvoted 1 times

youonebe 7 months, 1 week ago

Correct answer is D.

This is for NAT traffic analysis, so the focus is outbound.

VPC Flow Logs are published to CloudWatch Logs, not CloudTrail1. This immediately eliminates options A and C.

To determine if the traffic is unsolicited inbound connections:
We need to check if the private EC2 instance (starting with 203.0) initiated the connection to 198.51.100.2
If the source IP is from the VPC (203.0) and the destination is 198.51.100.2, this indicates the connection was initiated from inside the VPC
This would mean the ACCEPT traffic is a response to an outbound request, not unsolicited inbound traffic.

upvoted 2 times

tural_nasirov 7 months, 2 weeks ago

Selected Answer: B

The answer is B.

This is not about an IP but about a port. If the packet from outside to inside has the source port which is well known and the destination port dynamic, it means that the connection was initiated from inside, if the packet from outside to inside has a source port dynamic and destination port well known, it means that the traffic was originated from outside :)

upvoted 1 times

sashenka 7 months, 3 weeks ago

Selected Answer: D

Why Option B is Problematic:

// Example CloudWatch Logs Insights Query for Option B
fields @timestamp, sourceAddress, destinationAddress, action, bytes
| filter destinationAddress like "203.0"
| filter sourceAddress like "198.51.100.2"
| stats sum(bytes) by sourceAddress, destinationAddress

1. Incorrect Traffic Direction
- It looks for traffic where source = 198.51.100.2 (internet) and destination = 203.0.x.x (VPC)
This only shows successful inbound connections (ACCEPT)
It doesn't reveal whether these connections were solicited or unsolicited
2. Missing Context
- Doesn't show the initial outbound connection that would indicate a solicited response
- Cannot differentiate between legitimate responses and actual unsolicited connections
- Lacks the temporal relationship between outbound and inbound flows

upvoted 2 times

sashenka 7 months, 3 weeks ago

Better Approach (Option D)
sql
// Example CloudWatch Logs Insights Query for Option D
fields @timestamp, sourceAddress, destinationAddress, action, bytes
| filter sourceAddress like "203.0"
| filter destinationAddress like "198.51.100.2"
| stats sum(bytes) by sourceAddress, destinationAddress

This query would:
Show outbound traffic from VPC to the internet
Help establish if the private instance initiated communication
Allow correlation between outbound requests and inbound responses
Key Concept
With NAT gateway connections:
Legitimate traffic follows a request-response pattern
Outbound request must exist before inbound response
Looking only at inbound traffic (Option B) misses this crucial relationship
Therefore, Option D provides the necessary visibility to determine if the inbound connections were truly unsolicited by examining the outbound traffic first.

upvoted 1 times

sammyhaj 7 months, 4 weeks ago

D, we need to see if the internal origin was first used

upvoted 2 times

NirvanaSNM 11 months, 2 weeks ago

Selected Answer: B

destination address set as "like 203.0" and the source address set as "like 198.51.100.2"

upvoted 1 times

mns0173 12 months ago

Of course it is D. What useful info will you get from B? You need to check original request which in case of NAT is always EC2, not something in the internet.

upvoted 1 times

Helpnosense 1 year ago

Selected Answer: B

I vote B. Because the network traffic to check is unsolicited inbound connection. IT is initiated from the internet to internal EC2. The source is public IP address and the target is internal IP.

upvoted 1 times

higashikumi 1 year ago

Selected Answer: B

To determine whether the traffic represents unsolicited inbound connections from the internet, use the Amazon CloudWatch console. Select the log group that contains the NAT gateway's elastic network interface and the private instance's elastic network interface. Run a query to filter with the destination address set as "like 203.0" and the source address set as "like 198.51.100.2". This approach helps you analyze the VPC flow logs to identify if the inbound traffic to the private EC2 instance is expected return traffic or unsolicited. The stats command can be used to filter the sum of bytes transferred by the source address and the destination address, providing insight into the traffic patterns and ensuring network security.

upvoted 1 times

Vongolatt 1 year, 2 months ago

Selected Answer: D

the solution architect want to check if it's unsolicited traffic or not, so we need to check the if the request is sent by us. which means 198.51.100.2 should be the destination.

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B, CloudWatch & destination address 203.0

upvoted 2 times

ajeeshb 1 year, 3 months ago

Selected Answer: D

The question is "Solutions architect must determine whether the traffic represents unsolicited inbound connections from the internet". The NAT gateway does not allow any inbound traffic from an internet other than response to a traffic it sent out to internet which came from a VPC resource (eg, EC2). So to find out if the inbound traffic to NAT Gateway from internet IP 198.51.100.2 is unsolicit or not, check the vpc flowlog to see if there was an original request from source IP 203.0 to destination 198.51.100.2. This is what option D says.

upvoted 3 times

Question #91 *Topic 1*

A company consists or two separate business units. Each business unit has its own AWS account within a single organization in AWS Organizations. The business units regularly share sensitive documents with each other. To facilitate sharing, the company created an Amazon S3 bucket in each account and configured low-way replication between the S3 buckets. The S3 buckets have millions of objects.

Recently, a security audit identified that neither S3 bucket has encryption at rest enabled. Company policy requires that all documents must be stored with encryption at rest. The company wants to implement server-side encryption with Amazon S3 managed encryption keys (SSE-S3).

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

A. Turn on SSE-S3 on both S3 buckets. Use S3 Batch Operations to copy and encrypt the objects in the same location.

B. Create an AWS Key Management Service (AWS KMS) key in each account. Turn on server-side encryption with AWS KMS keys (SSE-KMS) on each S3 bucket by using the corresponding KMS key in that AWS account. Encrypt the existing objects by using an S3 copy command in the AWS CLI.

C. Turn on SSE-S3 on both S3 buckets. Encrypt the existing objects by using an S3 copy command in the AWS CLI.

D. Create an AWS Key Management Service, (AWS KMS) key in each account. Turn on server-side encryption with AWS KMS keys (SSE-KMS) on each S3 bucket by using the corresponding KMS key in that AWS account. Use S3 Batch Operations to copy the objects into the same location.

Suggested Answer: *C*

Community vote distribution

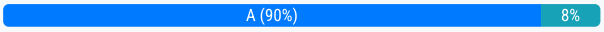

testingaws123 Highly Voted 2 years, 3 months ago

Selected Answer: A

Answer is A
Keyword is "The S3 buckets have millions of objects"
If there are million of objects then you should use Batch operations.
https://aws.amazon.com/blogs/storage/encrypting-objects-with-amazon-s3-batch-operations/
upvoted 27 times

forceli 2 years, 3 months ago

good point, changing my answer to A
upvoted 1 times

mnsait Most Recent 7 months ago

This is outdated now.
"Amazon S3 now applies server-side encryption with Amazon S3 managed keys (SSE-S3) as the base level of encryption for every bucket in Amazon S3. Starting January 5, 2023, all new object uploads to Amazon S3 are automatically encrypted at no additional cost and with no impact on performance."

https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html
upvoted 1 times

nimbus_00 8 months, 3 weeks ago

Selected Answer: A

S3 Batch Operations can be used to efficiently apply changes to a large number of objects in a bucket, including copying and encrypting them in place. This is ideal for retroactively encrypting millions of existing objects without needing to manually handle them one by one.
upvoted 1 times

ajeeshb 1 year, 3 months ago

I understand S3 Batch operations is required. But why no one is choosing SSE-KMS?
upvoted 1 times

StevePace 1 year, 3 months ago

Because the question states the company wants to use SSE-S3, nowhere does it mention SSE-KMS
upvoted 3 times

TonytheTiger 1 year, 3 months ago

To encrypt your existing unencrypted Amazon S3 objects, you can use Amazon S3 Batch Operations. You provide S3 Batch Operations with a list of objects to operate on, and Batch Operations calls the respective API to perform the specified operation. You can use the Batch Operations Copy operation to copy existing unencrypted objects and write them back to the same bucket as encrypted objects. A single Batch Operations job can perform the specified operation on billions of objects. https://docs.aws.amazon.com/AmazonS3/latest/userguide/bucket-encryption.html

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: A

A = correct (see https://aws.amazon.com/blogs/storage/encrypting-objects-with-amazon-s3-batch-operations/)

B = KMS is for SSE-KMS not for the requested SSE-S3

C = CLI is less efficient than S3 Batch

D = see answer B

upvoted 4 times

career360guru 1 year, 6 months ago

Selected Answer: A

A is the right answer

upvoted 1 times

jainparag1 1 year, 7 months ago

Selected Answer: A

Correct answer should be A. But this question seem too old to be true now since SSE-S3 based encryption is by default enabled and can't be disabled (you can change however) since Jan 2023.

upvoted 4 times

covabix879 1 year, 9 months ago

Selected Answer: D

Since SSE-S3 does not support cross-account replication, answer should be D

upvoted 2 times

deivid83 1 year, 9 months ago

In a cross-account scenario, where the source and destination buckets are owned by different AWS accounts, you can use a KMS key to encrypt object replicas. However, the KMS key owner must grant the source bucket owner permission to use the KMS key.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-config-for-kms-objects.html#replication-kms-cross-acct-scenario

S3 Batch operation:
https://aws.amazon.com/blogs/storage/encrypting-objects-with-amazon-s3-batch-operations/

upvoted 3 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: A

S3 Batch operation is the MOST operationally efficient way for millions objects

upvoted 1 times

sachstarinfoaws 1 year, 11 months ago

Selected Answer: A

Answer is A

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A more efficient

upvoted 1 times

Maria2023 2 years ago

Selected Answer: A

I vote for A. Batch operations is better for such a high number of objects

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: A

https://aws.amazon.com/blogs/storage/encrypting-objects-with-amazon-s3-batch-operations/
The launch of S3 default encryption feature automate the wok of encrypting new objects, and you asked for similar, straightforward ways to encrypt existing objects in your buckets. While tools and scripts exist to do this work, each one requires some development work to set up. S3 batch operations gives you a solution for encrypting large number of archived files.
This can also be done by CLI, Option C, however, the same article refers to Batch Operations in case you have a large bucket with millions of objects.
https://aws.amazon.com/blogs/storage/encrypting-existing-amazon-s3-objects-with-the-aws-cli/
Option A should be the most efficient, even though it has more operational cost to implement but the question is the about efficiency, it would take to much time to complete this using CLI (Option C).

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: A

A is much more efficient

upvoted 1 times

forceli 2 years, 3 months ago

Selected Answer: C

A and C seems to be correct but using batch requires more steps.
https://aws.amazon.com/blogs/storage/encrypting-existing-amazon-s3-objects-with-the-aws-cli/

upvoted 1 times

Question #92 *Topic 1*

A company is running an application in the AWS Cloud. The application collects and stores a large amount of unstructured data in an Amazon S3 bucket. The S3 bucket contains several terabytes of data and uses the S3 Standard storage class. The data increases in size by several gigabytes every day.

The company needs to query and analyze the data. The company does not access data that is more than 1 year old. However, the company must retain all the data indefinitely for compliance reasons.

Which solution will meet these requirements MOST cost-effectively?

A. Use S3 Select to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

B. Use Amazon Redshift Spectrum to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old 10 S3 Glacier Deep Archive.

C. Use an AWS Glue Data Catalog and Amazon Athena to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

D. Use Amazon Redshift Spectrum to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Intelligent-Tiering.

Suggested Answer: *A*

Community vote distribution

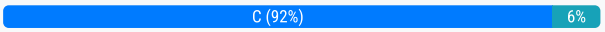

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct answer is C. Use an AWS Glue Data Catalog and Amazon Athena to query the data. Create an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive.

This solution allows you to use Amazon Athena and the AWS Glue Data Catalog to query and analyze the data in an S3 bucket. Amazon Athena is a serverless, interactive query service that allows you to analyze data in S3 using SQL. The AWS Glue Data Catalog is a managed metadata repository that can be used to store and retrieve table definitions for data stored in S3. Together, these services can provide a cost-effective way to query and analyze large amounts of unstructured data. Additionally, by using an S3 Lifecycle policy to transition data that is more than 1 year old to S3 Glacier Deep Archive, you can retain the data indefinitely for compliance reasons while also reducing storage costs.

upvoted 21 times

masetromain 2 years, 5 months ago

The other options are not correct because:

A. Using S3 Select is good for filtering data in S3, but it may not be a suitable solution for querying and analyzing large amounts of data.

B. Amazon Redshift Spectrum can be used to query data stored in S3, but it may not be as cost-effective as using Amazon Athena for querying unstructured data

D. Using Amazon Redshift Spectrum with S3 Intelligent-Tiering could be a good solution, but S3 Intelligent-Tiering is designed to optimize storage costs based on access patterns and it would not be the best solution for compliance reasons as S3 Intelligent-Tiering will move data to other storage classes according to access patterns.

upvoted 9 times

Japanese1 1 year, 7 months ago

This is a nonsense explanation.

In the first place, Redshift cannot handle unstructured data.

upvoted 4 times

dankositzke 1 year, 4 months ago

Amazon Redshift is designed for structured data. However, Amazon Redshift Spectrum enables you to run queries against exabytes of unstructured data in Amazon S3, with no loading or ETL required.

upvoted 3 times

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: C

Generally, unstructured data should be converted structured data before querying them. AWS Glue can do that.
https://docs.aws.amazon.com/glue/latest/dg/schema-relationalize.html
https://docs.aws.amazon.com/athena/latest/ug/glue-athena.html
upvoted 7 times

GabrielShiao Most Recent 8 months, 3 weeks ago

Selected Answer: C

B, C seem both acceptable. The reason C is selected is because redshift spectrum need Glue Data Catalog as well which is not mentioned there.
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, aws glue + amazon athena
upvoted 1 times

AimarLeo 1 year, 5 months ago

Many comments were not convincing of not using Redshift Spectrum.. the only reason I see it to exclude that option is a Redshift Spectrum MUST have a Redshift Cluster available to start the query to S3..
upvoted 1 times

djeong95 1 year, 3 months ago

This question is actually pretty difficult since both Redshift Spectrum and AWS Glue + Athena could query unstructured data. Redshift Spectrum and Athena actually cost about the same per TB. However, with Athena, you could lower the cost by compressing the data. Glue doesn't seem to cost that much either.

https://aws.amazon.com/redshift/pricing/
https://aws.amazon.com/athena/pricing/
https://aws.amazon.com/glue/pricing/
upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

A = S3 Select good for filtering an retrieve subset of data, not enough to analyze
B = need a Redshift instance that is expensive
C = correct (Glue Data Catalog can help putting some structure to data and Athena is good for both query and analytics, transition to Deep Archive after 1 year)
D = see answer B + Intelligent-Tiering not the best option here
upvoted 2 times

nzin4x 1 year, 5 months ago

redshift spectrum vs athena: https://www.upsolver.com/blog/aws-serverless-redshift-spectrum-athena

Both are good solutions to query s3 data. However, redshift spectrum is useful for joining S3 data with other data in Redshift, and if the data is only in S3, it would be preferable to choose athena.
upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

C is the right answer as Data needs to be queried and Analyzed.
upvoted 2 times

subbupro 1 year, 6 months ago

Athena and aws glue is more cost , so better go with A . and what is the purpose for aws glue here. AWS glue is for ETL purpose unnecessary
upvoted 1 times

Andy16240 1 year, 7 months ago

C correct: S3 copy command in AWS CLI is less operational processes than the batch operation.
upvoted 1 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: C

In this particular scenario, using Amazon Athena and AWS Glue Data Catalog might be a better fit due to the large amount of data stored in S3 buckets and growing every day. Athena can query data across an entire S3 bucket or across multiple buckets, which is useful when parsing multiple files and large amounts of data.
upvoted 2 times

chico2023 1 year, 10 months ago

Selected Answer: C

Answer: C

Criminally tricky question. S3 Select does the same thing as Athena but there are some differences. The key here is "...a large amount of unstructured data..."
If wasn't this, S3 Select hands down.
upvoted 3 times

chico2023 1 year, 10 months ago

Using an Olabiba to explain the differences between the two:

1. Query Capability: Amazon Athena is a fully managed interactive query service that allows you to run SQL queries directly on your data in S3. It supports complex queries, joins, aggregations, and even nested data structures. Athena is designed for ad-hoc querying and analysis of large datasets.

On the other hand, S3 Select is a feature of Amazon S3 that allows you to retrieve a subset of data from an object using SQL expressions. It is primarily used for selective retrieval of specific data within an object, rather than running complex queries across multiple objects.
upvoted 2 times

chico2023 1 year, 10 months ago

2. Data Format: Amazon Athena supports various data formats such as CSV, JSON, Parquet, Avro, and more. It can automatically infer the schema of your data or you can provide a schema explicitly. Athena can handle structured, semi-structured, and unstructured data.

S3 Select, on the other hand, is limited to querying CSV, JSON, and Parquet files. It requires the data to be in a specific format and does not support nested data structures.
upvoted 2 times

chico2023 1 year, 10 months ago

3. Performance: Amazon Athena is optimized for running queries on large datasets and can parallelize the query execution across multiple nodes. It automatically scales resources based on the query complexity and data size, providing fast and efficient query performance.

S3 Select, on the other hand, is designed for retrieving a subset of data from an object. It can significantly reduce the amount of data transferred over the network and improve query performance by only retrieving the necessary data.

4. Cost: Both Amazon Athena and S3 Select have different pricing models. Amazon Athena charges based on the amount of data scanned by your queries, while S3 Select charges based on the amount of data selected and returned by your queries. The cost will depend on the size of your data and the complexity of your queries.
upvoted 3 times

Jonalb 1 year, 11 months ago

Selected Answer: C

its a C , true question!
upvoted 1 times

NikkyDicky 1 year, 12 months ago

C for sure
upvoted 1 times

johnballs221 2 years, 1 month ago

Selected Answer: B

redshift spectrum can run sql queries directly on s3
upvoted 1 times

rxhan 2 years ago

Not the best for cost.
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

C is the best choice for unstructured data

upvoted 3 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

S3 select only to select few parts of the data and here its lot of unstructured data. So A is wrong. Use Athena console to create Glue crawler as referred here -

https://docs.aws.amazon.com/athena/latest/ug/data-sources-glue.html

upvoted 4 times

Question #93 *Topic 1*

A video processing company wants to build a machine learning (ML) model by using 600 TB of compressed data that is stored as thousands of files in the company's on-premises network attached storage system. The company does not have the necessary compute resources on premises for ML experiments and wants to use AWS.

The company needs to complete the data transfer to AWS within 3 weeks. The data transfer will be a one-time transfer. The data must be encrypted in transit. The measured upload speed of the company's internet connection is 100 Mbps. and multiple departments share the connection.

Which solution will meet these requirements MOST cost-effectively?

A. Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console. Configure the devices with a destination S3 bucket. Copy the data to the devices. Ship the devices back to AWS.

B. Set up a 10 Gbps AWS Direct Connect connection between the company location and the nearest AWS Region. Transfer the data over a VPN connection into the Region to store the data in Amazon S3.

C. Create a VPN connection between the on-premises network attached storage and the nearest AWS Region. Transfer the data over the VPN connection.

D. Deploy an AWS Storage Gateway file gateway on premises. Configure the file gateway with a destination S3 bucket. Copy the data to the file gateway.

Suggested Answer: *A*

Community vote distribution

A (100%)

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: A

The correct answer is A. Order several AWS Snowball Edge Storage Optimized devices by using the AWS Management Console. Configure the devices with a destination S3 bucket. Copy the data to the devices. Ship the devices back to AWS.

This option will meet the requirements to complete the data transfer within 3 weeks, as the Snowball Edge devices can transfer large amounts of data quickly and securely. The data will be encrypted in transit and at rest. The company's internet connection speed is not a bottleneck as the data transfer will happen on the devices and not over the internet.

upvoted 11 times

masetromain 1 year, 11 months ago

Option B is not a cost-effective solution, as setting up and maintaining a 10 Gbps Direct Connect connection can be quite expensive, especially if it's only needed for a one-time data transfer.

Option C is not a cost-effective solution, as creating a VPN connection between the on-premises storage and the nearest AWS region would require significant networking configuration and maintenance, and would likely be more expensive than using Snowball Edge devices.

Option D is not a cost-effective solution, as deploying an AWS Storage Gateway file gateway on premises would require additional hardware and ongoing maintenance costs, and may not be necessary for a one-time data transfer.

upvoted 3 times

ninomfr64 Most Recent 11 months, 2 weeks ago

Selected Answer: A

A = correct

B = takes a month or more to setup DX

C = this would take more than 3 weeks for transferring data

D = this would take more than 3 weeks for transferring data

upvoted 2 times

career360guru 1 year ago

Selected Answer: A

Option A

upvoted 1 times

yorkicurke 1 year, 2 months ago

Selected Answer: A

wish all the questions were like this. happy days :)

upvoted 1 times

xplusfb 1 year, 4 months ago

Selected Answer: A

as we know snowball storage optimized NVMe up to 210 TB <3 A is the best and easy answer

upvoted 4 times

xplusfb 1 year, 4 months ago

like several sorry for any confision :)

upvoted 1 times

chikorita 1 year, 4 months ago

several thanks too :)

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

A - basic snowball use case

upvoted 1 times

Maria2023 1 year, 6 months ago

Selected Answer: A

Given the deadline (3 weeks) and the amount of data I would use Snowball Edge

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: A

A obviously

upvoted 3 times

God_Is_Love 1 year, 9 months ago

Selected Answer: A

Around 8 devices and snowball (actually a Rectangular box)

Snowball Edge Storage Optimized device is equipped with up to 80 terabytes (TB) of storage capacity, as well as 40 vCPUs and 80 GB of memory for running compute-intensive applications. It also includes an optional GPU for accelerated computing workloads.

Built-in security features such as tamper-resistant enclosures, an E Ink shipping label, and 256-bit encryption for data at rest and in transit.

upvoted 4 times

zozza2023 1 year, 11 months ago

Selected Answer: A

3 weeks + cost effective ==> Snowball Edge Storage

upvoted 1 times

Question #94 *Topic 1*

A company has migrated Its forms-processing application to AWS. When users interact with the application, they upload scanned forms as files through a web application. A database stores user metadata and references to files that are stored in Amazon S3. The web application runs on Amazon EC2 instances and an Amazon RDS for PostgreSQL database.

When forms are uploaded, the application sends notifications to a team through Amazon Simple Notification Service (Amazon SNS). A team member then logs in and processes each form. The team member performs data validation on the form and extracts relevant data before entering the information into another system that uses an API.

A solutions architect needs to automate the manual processing of the forms. The solution must provide accurate form extraction. minimize time to market, and minimize tong-term operational overhead.

Which solution will meet these requirements?

A. Develop custom libraries to perform optical character recognition (OCR) on the forms. Deploy the libraries to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster as an application tier. Use this tier to process the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data into an Amazon DynamoDB table. Submit the data to the target system's APL. Host the new application tier on EC2 instances.

B. Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use artificial intelligence and machine learning (AI/ML) models that are trained and hosted on an EC2 instance to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API.

C. Host a new application tier on EC2 instances. Use this tier to call endpoints that host artificial intelligence and machine teaming (AI/ML) models that are trained and hosted in Amazon SageMaker to perform optical character recognition (OCR) on the forms. Store the output in Amazon ElastiCache. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API.

D. Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API.

Suggested Answer: *D*

Community vote distribution

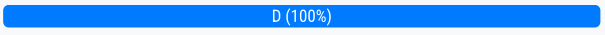

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: D

The correct answer is D. Extend the system with an application tier that uses AWS Step Functions and AWS Lambda. Configure this tier to use Amazon Textract and Amazon Comprehend to perform optical character recognition (OCR) on the forms when forms are uploaded. Store the output in Amazon S3. Parse this output by extracting the data that is required within the application tier. Submit the data to the target system's API.

This solution meets the requirements of accurate form extraction, minimal time to market, and minimal long-term operational overhead. Amazon Textract and Amazon Comprehend are fully managed and serverless services that can perform OCR and extract relevant data from the forms, which eliminates the need to develop custom libraries or train and host models. Using AWS Step Functions and Lambda allows for easy automation of the process and the ability to scale as needed.

upvoted 15 times

masetromain 1 year, 11 months ago

Option A:

This option would require significant development and maintenance effort and would not take advantage of fully managed services, resulting in increased operational overhead.

Option B:

This option is similar to option A in that it would require significant development and maintenance effort to train and host the models, and would not take advantage of fully managed services resulting in increased operational overhead.

Option C:
This option is similar to option B in that it would require significant development and maintenance effort to train and host the models, and would not take advantage of fully managed services resulting in increased operational overhead.
upvoted 3 times

gofavad926 Most Recent 9 months, 2 weeks ago

Selected Answer: D

D. This solution meets the requirements of accurate form extraction, minimal time to market, and minimal long-term operational overhead
upvoted 1 times

career360guru 1 year ago

Selected Answer: D

Option D
upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: D

D - basic use case for textract
upvoted 1 times

Maria2023 1 year, 6 months ago

Selected Answer: D

An easy one - if AWS has a service for something - do not reinvent the wheel - use Textract and Comprehend
upvoted 2 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: D

D : Managed AWS Services
upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: D

Amazon Textract..
upvoted 1 times

God_Is_Love 1 year, 9 months ago

Selected Answer: D

Textract can analyze different types of documents such as forms, invoices, receipts, and tables, and can extract information such as text, tables, and key-value pairs.

Comprehend provides a set of APIs that can be used to analyze text data in real-time. The service can identify the language of the text, extract entities such as people, organizations, and locations, and detect the sentiment expressed in the text. It can also extract key phrases that summarize the meaning of the text, and can classify the text into predefined categories.
upvoted 1 times

sambb 1 year, 10 months ago

Selected Answer: D

D : Managed AWS Services
upvoted 1 times

Question #95 *Topic 1*

A company is refactoring its on-premises order-processing platform in the AWS Cloud. The platform includes a web front end that is hosted on a fleet of VMs, RabbitMQ to connect the front end to the backend, and a Kubernetes cluster to run a containerized backend system to process the orders. The company does not want to make any major changes to the application.

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

A. Create an AMI of the web server VM. Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer. Set up Amazon MQ to replace the on-premises messaging queue. Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend.

B. Create a custom AWS Lambda runtime to mimic the web server environment. Create an Amazon API Gateway API to replace the front-end web servers. Set up Amazon MQ to replace the on-premises messaging queue. Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend.

C. Create an AMI of the web server VM. Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer. Set up Amazon MQ to replace the on-premises messaging queue. Install Kubernetes on a fleet of different EC2 instances to host the order-processing backend.

D. Create an AMI of the web server VM. Create an Amazon EC2 Auto Scaling group that uses the AMI and an Application Load Balancer. Set up an Amazon Simple Queue Service (Amazon SQS) queue to replace the on-premises messaging queue. Configure Amazon Elastic Kubernetes Service (Amazon EKS) to host the order-processing backend.

Suggested Answer: *A*

Community vote distribution

A (93%) 7%

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

Option A is the correct answer. In this solution, the company creates an Amazon Machine Image (AMI) of the web server VM, which can be used to launch EC2 instances that are identical to the on-premises web servers. The company then creates an EC2 Auto Scaling group that uses the AMI and an Application Load Balancer (ALB) to provide automatic scaling and high availability for the web front end. The company also replaces the on-premises messaging queue (RabbitMQ) with Amazon MQ, which is a managed message broker service that is fully compatible with RabbitMQ. Finally, the company uses Amazon Elastic Kubernetes Service (EKS) to host the order-processing backend, which allows them to run their existing Kubernetes cluster in the AWS cloud without making any major changes to the application. This approach allows the company to lift and shift their existing platform with minimal operational overhead.

upvoted 20 times

pk0619 6 months, 1 week ago

AMI is an AWS EC2 specific, I am confused on how to create an AMI of an on-premise VM and launch instance from it ?

upvoted 1 times

pk0619 6 months, 1 week ago

Looking back, it seems the intent might have been to use VM Import/Export or the AWS Application Migration Service (MGN) to create an AMI. However, it is a significant oversimplification to skip those critical steps and simply state "create an AMI," as this assumes the process is straightforward without addressing the necessary prerequisites and tools.

upvoted 1 times

masetromain 2 years, 5 months ago

Option B, using a custom AWS Lambda runtime and Amazon API Gateway, would require significant changes to the application and may not be compatible with the current codebase.

Option C, installing Kubernetes on a fleet of different EC2 instances, would also require significant changes to the application and may not be compatible with the current codebase.

Option D, using Amazon Simple Queue Service (Amazon SQS) instead of Amazon MQ, would not provide the same level of messaging capabilities as Amazon MQ and may not be sufficient for the needs of the order-processing platform.

upvoted 4 times

sambb 2 years, 3 months ago

Your justification for option C is wrong.

Option C is valid, as Kubernetes on EC2 is very similar as the existing Kubernetes environment on-premises. But EKS is a safe bet and reduces operational overhead, while keeping the same API as previously. Hence, A is a better choice.

upvoted 10 times

AimarLeo Highly Voted 1 year, 5 months ago

AWS exams got more 'sarcastic' with the ways of formulating questions.. E.g here: _'A company is refactoring its on-premises order-processing platform in the AWS Cloud'

BUT '

The company does not want to make any major changes to the application.

Replatforming and Rehosting is not real refactoring.. but the closest answer as an architect with least operational overhead is A obvisouly.. aws questions sometimes can be ultra vague

upvoted 6 times

madeesha Most Recent 1 year ago

Selected Answer: A

answer is A

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A, is the only option to don't involve a rearchitectured solution

upvoted 1 times

jpa8300 1 year, 5 months ago

Selected Answer: A

A better explanation to choose between option A and D is that Amazon MQ respondes to the requirement of not changing the app, because it accepts the same protocol as RabbitMQ (Supports AMQP, MQTT, STOMP, OpenWire, and JMS) while SQS has its own API, so it would need more changes to the app.

upvoted 3 times

career360guru 1 year, 6 months ago

Selected Answer: A

Option A

upvoted 1 times

Mikado211 1 year, 7 months ago

Selected Answer: A

a bunch of keywords for this migration here :

Kubernetes == EKS

RabbitMQ == Amazon MQ

A fleet of VM == AMI + ec2 instances

The answer A proposes all thoses points, so it's perfect here.

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A no doubt

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: A

A is the best choice.

upvoted 1 times

Musk 2 years, 4 months ago

Selected Answer: B

Option A is re-hosting or mybe re-platforming. The question says the purpose is re-factoring, then it's B.

upvoted 2 times

c73bf38 2 years, 4 months ago

It says the company does not want to make changes to the application in the problem statement. B would require significant code changes to the application.

upvoted 6 times

Question #96 *Topic 1*

A solutions architect needs to implement a client-side encryption mechanism for objects that will be stored in a new Amazon S3 bucket. The solutions architect created a CMK that is stored in AWS Key Management Service (AWS KMS) for this purpose.

The solutions architect created the following IAM policy and attached it to an IAM role:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "DownloadUpload",
            "Action": [
                "s3:GetObject",
                "s3:GetObjectVersion",
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:s3:::BucketName/*"
        },
        {
            "Sid": "KMSAccess",
            "Action": [
                "kms:Decrypt",
                "kms:Encrypt"
            ],
            "Effect": "Allow",
            "Resource": "arn:aws:kms:Region:Account:key/Key ID"
        }
    ]
}
```

During tests, the solutions architect was able to successfully get existing test objects in the S3 bucket. However, attempts to upload a new object resulted in an error message. The error message stated that the action was forbidden.

Which action must the solutions architect add to the IAM policy to meet all the requirements?

A. kms:GenerateDataKey

B. kms:GetKeyPolicy

C. kms:GetPublicKey

D. kms:Sign

Suggested Answer: *A*

Community vote distribution

A (100%)

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: A

A. kms:GenerateDataKey

The solutions architect needs to add the "kms:GenerateDataKey" action to the IAM policy in order to generate a data key for client-side encryption. Without this action, the IAM role does not have the necessary permissions to generate a data key, which causes the error message when attempting to upload a new object.

upvoted 15 times

masetromain 1 year, 11 months ago

The other options are not correct because they are not required for this use case. kms:GetKeyPolicy allows for the retrieval of the key policy for a CMK but it does not have any relation to client-side encryption of S3 objects, kms:GetPublicKey allows for the retrieval of the public key of a CMK, but it does not have any relation to client-side encryption of S3 objects and kms:Sign allows for signing a message using a CMK but it does not have any relation to client-side encryption of S3 objects.

upvoted 2 times

altonh Most Recent 5 months, 2 weeks ago

Selected Answer: A

The answers don't make sense. The requirement is that it will be client-side encryption, which means the object is already encrypted when sent to S3. S3 will not do any encryption, so S3 does not need to access the KMS key,

upvoted 1 times

ninomfr64 11 months, 2 weeks ago

Selected Answer: A

A = correct (you encrypt data with KMS Data Key and not KMS Key directly, unless data is < 4K)

B = getting the policy would allow to get the data key needed for encryption

C = client side encryption uses symmetric key not asymmetric keys

D = sign allows for signing messages, API calls, etc.

upvoted 3 times

career360guru 1 year ago

Selected Answer: A

Option A

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

A - need data key for client-side encr

upvoted 1 times

Jesuisleon 1 year, 7 months ago

I don't understand since it's client side encryption, it means both encryption and key and tools are maintained in client side before submitting to aws s3, why we need add kms:GenerateDatakey ? We don't need kms to do anything since it's client-side encryption all is done outside of aws.

upvoted 4 times

venvig 1 year, 4 months ago

When you want to do the client side encryption, your files are most likely above 4K in size. So, you would be performing envelope encryption.

For that, you need a data key.

You ask KMS to generate and give you the data key, supplying the kms CMK.

KMS would generate a new data key, encrypt it with the CMK and return you both the encrypted and plain data key. AWS would never retain the data key; they will immediately discard it.

You would now encrypt your data using the plain data key and immediately delete the plain data key (unencrypted). You store the encrypted data key that you got from KMS along with the encrypted data, which is then uploaded to s3. Note that AWS does NOT know about the data key at this point; only you know. KMS just holds the kms CMK that was used to encrypt the data key.

So, you need access to KMS to decrypt the data key before using that decrypted data key to unencrypt your data.

Similarly AWS cannot read your data, even though it has the KMS CMK and also the encrypted data key stored in s3.

This is why you need the generateDataKey permission. Hope this helps.

upvoted 11 times

venvig 1 year, 4 months ago

Of course the answer is A

upvoted 1 times

bcx 1 year, 6 months ago

Indeed, the question says client side encryption but the answer is all about S3-KMS.

upvoted 2 times

mfsec 1 year, 9 months ago

Selected Answer: A

A for sure

upvoted 1 times

Untamables 1 year, 11 months ago

Selected Answer: A

https://docs.aws.amazon.com/kms/latest/cryptographic-details/client-side-encryption.html

upvoted 3 times

masssa 1 year, 11 months ago

Selected Answer: A

I Vote A.

https://repost.aws/ja/knowledge-center/s3-large-file-encryption-kms-key

Adding kms:GenerateDataKey is necessary.

upvoted 1 times

Question #97 *Topic 1*

A company has developed a web application. The company is hosting the application on a group of Amazon EC2 instances behind an Application Load Balancer. The company wants to improve the security posture of the application and plans to use AWS WAF web ACLs. The solution must not adversely affect legitimate traffic to the application.

How should a solutions architect configure the web ACLs to meet these requirements?

A. Set the action of the web ACL rules to Count. Enable AWS WAF logging. Analyze the requests for false positives. Modify the rules to avoid any false positive. Over time, change the action of the web ACL rules from Count to Block.

B. Use only rate-based rules in the web ACLs, and set the throttle limit as high as possible. Temporarily block all requests that exceed the limit. Define nested rules to narrow the scope of the rate tracking.

C. Set the action of the web ACL rules to Block. Use only AWS managed rule groups in the web ACLs. Evaluate the rule groups by using Amazon CloudWatch metrics with AWS WAF sampled requests or AWS WAF logs.

D. Use only custom rule groups in the web ACLs, and set the action to Allow. Enable AWS WAF logging. Analyze the requests for false positives. Modify the rules to avoid any false positive. Over time, change the action of the web ACL rules from Allow to Block.

Suggested Answer: *A*

Community vote distribution

A (100%)

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: A

AWS WAF allows you to create web ACL (Access Control List) rules in "Count" mode, which allows you to monitor traffic without actually blocking it. In Count mode, AWS WAF counts the number of requests that match a particular rule, but doesn't take any action to block those requests.

Count mode can be useful in several ways:

Testing new rules: You can create new rules and test them in Count mode before enabling them to block traffic. This allows you to evaluate the effectiveness of your rules without risking false positives or false negatives.

Analyzing traffic: You can use Count mode to analyze traffic patterns and identify potential security threats. By monitoring the number of requests that match a particular rule, you can detect patterns that may indicate an attack or vulnerability.

Compliance reporting: Count mode can be used for compliance reporting, where you need to demonstrate that certain rules are being enforced. By counting the number of requests that match a rule, you can provide evidence that your security policies are being followed.

upvoted 22 times

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: A

https://www.examtopics.com/discussions/amazon/view/74273-exam-aws-certified-solutions-architect-professional-topic-1/

The correct answer is A. Set the action of the web ACL rules to Count. Enable AWS WAF logging. Analyze the requests for false positives. Modify the rules to avoid any false positive. Over time, change the action of the web ACL rules from Count to Block.

This approach allows for monitoring of the incoming traffic and its behavior before taking any action that can affect the legitimate traffic. By setting the action to count, the web ACL will only log the requests that match the conditions of the rules, but it will not block them. This way, the company can analyze the requests and check for any false positives. Once they identify and correct any false positives, they can gradually change the action of the web ACL rules from count to block, thus improving the security posture of the application without adversely affecting legitimate traffic.

upvoted 6 times

masetromain 1 year, 11 months ago

Option B is not correct because using only rate-based rules can lead to false positives and blocking of legitimate traffic. Option C is not correct because using only AWS managed rule groups can limit the flexibility and specificity of the web ACLs. Option D is not correct because using only custom rule groups with action set to allow can lead to security vulnerabilities.

upvoted 1 times

gofavad926 Most Recent 9 months, 2 weeks ago

Selected Answer: A

A, configure the rules on COUNT

upvoted 1 times

Explorer_30 1 year, 4 months ago

vote A

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

Its an A

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: A

A. Set the action of the web ACL rules to Count. Enable AWS WAF logging.

upvoted 1 times

Untamables 1 year, 11 months ago

Selected Answer: A

https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-testing.html

upvoted 1 times

Question #98 *Topic 1*

A company has an organization that has many AWS accounts in AWS Organizations. A solutions architect must improve how the company manages common security group rules for the AWS accounts in the organization.

The company has a common set of IP CIDR ranges in an allow list in each AWS account to allow access to and from the company's on-premises network. Developers within each account are responsible for adding new IP CIDR ranges to their security groups. The security team has its own AWS account. Currently, the security team notifies the owners of the other AWS accounts when changes are made to the allow list.

The solutions architect must design a solution that distributes the common set of CIDR ranges across all accounts.

Which solution meets these requirements with the LEAST amount of operational overhead?

A. Set up an Amazon Simple Notification Service (Amazon SNS) topic in the security team's AWS account. Deploy an AWS Lambda function in each AWS account. Configure the Lambda function to run every time an SNS topic receives a message. Configure the Lambda function to take an IP address as input and add it to a list of security groups in the account. Instruct the security team to distribute changes by publishing messages to its SNS topic.

B. Create new customer-managed prefix lists in each AWS account within the organization. Populate the prefix lists in each account with all internal CIDR ranges. Notify the owner of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups. Instruct the security team to share updates with each AWS account owner.

C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.

D. Create an IAM role in each account in the organization. Grant permissions to update security groups. Deploy an AWS Lambda function in the security team's AWS account. Configure the Lambda function to take a list of internal IP addresses as input, assume a role in each organization account, and add the list of IP addresses to the security groups in each account.

Suggested Answer: *C*

Community vote distribution

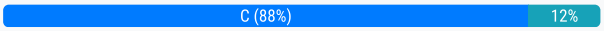

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: C

C. Create a new customer-managed prefix list in the security team's AWS account. Populate the customer-managed prefix list with all internal CIDR ranges. Share the customer-managed prefix list with the organization by using AWS Resource Access Manager. Notify the owner of each AWS account to allow the new customer-managed prefix list ID in their security groups.

This solution meets the requirements with the least amount of operational overhead as it requires the security team to create and maintain a single customer-managed prefix list, and share it with the organization using AWS Resource Access Manager. The owners of each AWS account are then responsible for allowing the prefix list in their security groups, which eliminates the need for the security team to manually notify each account owner when changes are made. This solution also eliminates the need for a separate AWS Lambda function in each account, reducing the overall complexity of the solution.

upvoted 11 times

masetromain 1 year, 11 months ago

Option A is not correct because it requires setting up an SNS topic in the security team's AWS account, and deploying an AWS Lambda function in each AWS account. This increases the operational overhead as it requires setting up and maintaining the SNS topic, and deploying and configuring the Lambda function in each account.

Option B is not correct because it requires creating new customer-managed prefix lists in each AWS account within the organization, which increases the operational overhead as it requires the security team to create and maintain multiple prefix lists.

Option D is not correct because it requires creating an IAM role in each account in the organization, which increases the operational overhead as it requires the security team to set up and maintain multiple roles. Additionally, it also deploys an AWS Lambda function in the security team's AWS account, which increases complexity and operational overhead.

upvoted 2 times

bur4an Highly Voted 1 year, 3 months ago

masetromain is ChatGPT and might have outdated answers since it doesnt know aws latest update to services

upvoted 8 times

AlbertC Most Recent 9 months, 1 week ago

Human cost is major overhead. I will go A. This is one time setup.

upvoted 1 times

StevePace 9 months, 2 weeks ago

Selected Answer: C

Centralised management and standard use case for prefix lists and RAM

upvoted 1 times

career360guru 1 year ago

Selected Answer: C

Option C

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

C - basic RAM use case

upvoted 1 times

bcx 1 year, 6 months ago

Selected Answer: C

Typical use case for RAM. It is the typical question that leads you to the solution without even finishing reading the question.

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: C

KEYWORD = AWS Resource Access Manager

Then C

upvoted 1 times

johnballs221 1 year, 7 months ago

Selected Answer: D

operational overhead

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: C

Prefix lists + RAM

upvoted 2 times

God_Is_Love 1 year, 9 months ago

Prefix lists + Resource Access Manager RAM is the solution.

upvoted 5 times

Musk 1 year, 11 months ago

Selected Answer: C

Clearly

upvoted 1 times

zozza2023 1 year, 11 months ago

Selected Answer: C

Create a new customer-managed prefix list in the security team's AWS account

upvoted 1 times

Untamables 1 year, 11 months ago

Selected Answer: C

https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html

upvoted 3 times

zhangyu20000 1 year, 11 months ago

C is correct. The prefix list is managed by security team and shared with other accounts. Other accounts can directly use it.

upvoted 1 times

masetromain 1 year, 11 months ago

Selected Answer: D

The correct answer is D.

Option D creates an IAM role in each account in the organization which grants permissions to update security groups. Then, it deploys an AWS Lambda function in the security team's AWS account, this lambda function is able to assume the IAM roles in each account and update the security groups with the new IP CIDR ranges. This solution allows the security team to easily distribute and update the common set of IP CIDR ranges across all accounts with minimal operational overhead.

Option A, uses an SNS topic, where the security team would need to notify all account owners every time an update is made to the allow list and would require the developers in each account to run a Lambda function which updates the security group. This solution would require a lot of manual work, and is not automated.

upvoted 2 times

masetromain 1 year, 11 months ago

Option B, requires the security team to notify the owners of each AWS account to allow the new customer-managed prefix list IDs in their accounts in their security groups, this solution would not provide a centralized control of the IP CIDR ranges and would require a lot of manual work.

Option C, uses a customer-managed prefix list in the security team's AWS account. But, it still requires the owners of each account to allow the new customer-managed prefix list ID in their security groups, this solution would not provide a centralized control of the IP CIDR ranges and would require a lot of manual work.

upvoted 1 times

God_Is_Love 1 year, 9 months ago

Create an IAM role in each account in the organization. this does not add up to operational overhead right.

upvoted 1 times

BabaP 1 year, 6 months ago

It's ChatGPT talking

upvoted 1 times

Question #99 *Topic 1*

A company has introduced a new policy that allows employees to work remotely from their homes if they connect by using a VPN. The company is hosting internal applications with VPCs in multiple AWS accounts. Currently, the applications are accessible from the company's on-premises office network through an AWS Site-to-Site VPN connection. The VPC in the company's main AWS account has peering connections established with VPCs in other AWS accounts.

A solutions architect must design a scalable AWS Client VPN solution for employees to use while they work from home.

What is the MOST cost-effective solution that meets these requirements?

A. Create a Client VPN endpoint in each AWS account. Configure required routing that allows access to internal applications.

B. Create a Client VPN endpoint in the main AWS account. Configure required routing that allows access to internal applications.

C. Create a Client VPN endpoint in the main AWS account. Provision a transit gateway that is connected to each AWS account. Configure required routing that allows access to internal applications.

D. Create a Client VPN endpoint in the main AWS account. Establish connectivity between the Client VPN endpoint and the AWS Site-to-Site VPN.

Suggested Answer: *B*

Community vote distribution

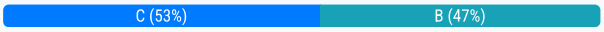

hexie Highly Voted 1 year, 12 months ago

Selected Answer: C

C.

Have you guys worked in a place where the configuration of B works?

The question clearly ask to design something scalable, and on C, the Transit Gateway serves as a network transit hub, allowing VPN connections to access resources across multiple VPCs in different AWS accounts.

VPC peering connections do not support transitive peering relationships, which means that if a user is connected to one VPC via AWS Client VPN, they cannot access resources in another VPC that's connected via a peering connection.

upvoted 36 times

Impromptu 1 year, 6 months ago

The question asks a scalable Client VPN solution (i.e. no openvpn on an EC2 instance or something like that), and asks for the most cost-effective. So AWS Client VPN is the scalable option. Reusing the current VPC peering is the most cost-effective compared to the far more expensive transit gateway solution.

I do agree that the peering does not support transitive peering. But for AWS Client VPN you get an ENI in the main account VPC and using the ENI you can access the VPCs over the VPC peering. So that does really work (in contrast to the Site-To-Site VPN):

https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/scenario-peered.html

upvoted 13 times

Jassybanga 10 months ago

Most cost effective - Transit gateway option is more costlier then B

upvoted 3 times

artazar 3 months, 3 weeks ago

Direct link from the docs for the scenario:

https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/how-it-works.html#scenario-peered

Transitive peering is VPC A <-peer-> VPC B <-peer-> VPC C ---> here VPC A cannot communicate to VPC C. But Client VPN is not a peering connection.

upvoted 1 times

vn_thanhtung 1 year, 10 months ago

The VPC in the company's main AWS account has peering connections established with VPCs in other AWS accounts => no need transit gw

upvoted 13 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

https://www.examtopics.com/discussions/amazon/view/80782-exam-aws-certified-solutions-architect-professional-topic-1/

B. Create a Client VPN endpoint in the main AWS account. Configure required routing that allows access to internal applications is the MOST cost-effective solution that meets these requirements. This solution allows employees to connect to the main AWS account using a Client VPN endpoint, and then use peering connections established with other AWS accounts to access the internal applications. This eliminates the need for additional Client VPN endpoints in each AWS account, reducing costs.

Option A, creating a Client VPN endpoint in each AWS account, would be more expensive as it would require multiple endpoints.

Option C, creating a transit gateway, would also add unnecessary costs.

Option D, connecting the Client VPN endpoint to the Site-to-Site VPN, may not provide a scalable solution for remote employees.

upvoted 24 times

Kaps443 Most Recent 2 weeks, 6 days ago

Selected Answer: C

Option C is the BEST solution: it provides a centralized, scalable, and cost-effective VPN access architecture using AWS Client VPN + Transit Gateway to allow secure access across multiple AWS accounts and VPCs.

upvoted 1 times

jimee11 1 month ago

Selected Answer: C

Transitive will cause issues trying to connect too her VPCs.

upvoted 1 times

eesa 1 month, 2 weeks ago

Selected Answer: B

B. Crear un Client VPN endpoint en la cuenta principal y configurar enrutamiento

✅ Muy buena opción:

Un solo Client VPN compartido para todos los usuarios.

Como la cuenta principal ya tiene peering con otras cuentas/VPCs, puedes simplemente rutar el tráfico hacia ellas desde el Client VPN.

Esto es sencillo, escalable y muy rentable.

upvoted 1 times

BennyMao 3 months, 3 weeks ago

Selected Answer: C

This provides a scalable and centralized routing solution to connect VPCs across multiple AWS accounts.

upvoted 1 times

Liliwood 5 months, 2 weeks ago

Selected Answer: B

Option B is the most cost-effective solution as it only requires creating a single Client VPN endpoint in the main AWS account and configuring the required routing to access the internal applications across the VPC peering connections.

Option C would involve additional costs for provisioning a transit gateway and connecting it to each AWS account, which is not necessary in this scenario since the VPCs are already peered.

upvoted 1 times

henrikhmkhitaryan59 7 months ago

Selected Answer: B

Option B is the MOST cost-effective solution that meets the requirements.

upvoted 2 times

Hibiki761 7 months, 1 week ago

Selected Answer: B

VPC peering is enough

upvoted 1 times

0b43291 7 months, 1 week ago

Selected Answer: B

By choosing option B, you can provide a scalable and cost-effective solution for remote employees to access internal applications hosted in multiple AWS accounts, while leveraging the existing VPC peering connections and minimizing the number of AWS resources required.

The other options are either more complex, less cost-effective, or introduce unnecessary components:
A. Creating a Client VPN endpoint in each AWS account would be more expensive and harder to manage, as you would need to configure and maintain multiple endpoints.
C. Provisioning a Transit Gateway in addition to the Client VPN endpoint would introduce an additional service and associated costs, which may not be necessary if the existing VPC peering connections are sufficient.
D. Establishing connectivity between the Client VPN endpoint and the AWS Site-to-Site VPN would introduce unnecessary complexity, as the Site-to-Site VPN is intended for connecting the on-premises office network, not individual remote employees.

upvoted 1 times

youonebe 7 months, 1 week ago

answer is B, should take advantage of existing VPC peering connections which works with current network topology

upvoted 1 times

Halliphax 7 months, 3 weeks ago

Selected Answer: B

B.
It asks for a scalable solution and it has to be cost effective. Adding Transit Gateway is not cost effective and also not required as the main AWS account has peering connections to VPCs in other accounts already.

upvoted 1 times

sammyhaj 7 months, 4 weeks ago

Selected Answer: B

No tgw needed

upvoted 2 times

Johnoppong101 9 months, 2 weeks ago

Selected Answer: C

Always find possible solutions first. Then look for cost effective. A cost effective option that does not solve the requirements is by default the most expensive option.

Requirement: most scalable option

upvoted 1 times

amministrazione 10 months ago

B. Create a Client VPN endpoint in the main AWS account. Configure required routing that allows access to internal applications.

upvoted 1 times

Syre 11 months ago

Selected Answer: B

C introduces additional costs.

upvoted 1 times

zolthar_z 11 months, 1 week ago

Selected Answer: B

Answer is B, right now you have VPC Peering from main VPC to others account VPC, you can re-use that configuration, also transit-gateway has a cost based on connections and traffic and the solution must be MOST cost-effective

upvoted 2 times

Question #100 *Topic 1*

A company is running an application in the AWS Cloud. Recent application metrics show inconsistent response times and a significant increase in error rates. Calls to third-party services are causing the delays. Currently, the application calls third-party services synchronously by directly invoking an AWS Lambda function.

A solutions architect needs to decouple the third-party service calls and ensure that all the calls are eventually completed.

Which solution will meet these requirements?

A. Use an Amazon Simple Queue Service (Amazon SQS) queue to store events and invoke the Lambda function.

B. Use an AWS Step Functions state machine to pass events to the Lambda function.

C. Use an Amazon EventBridge rule to pass events to the Lambda function.

D. Use an Amazon Simple Notification Service (Amazon SNS) topic to store events and Invoke the Lambda function.

Suggested Answer: *A*

Community vote distribution

A (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

The correct answer is A. Using an Amazon Simple Queue Service (SQS) queue to store events and invoke the Lambda function is a good solution to decouple the third-party service calls and ensure that all the calls are eventually completed. SQS is a fully managed, reliable, and highly scalable message queuing service that allows applications to send, store, and receive messages between distributed components. By sending the third-party service calls to an SQS queue, it allows the application to continue processing without waiting for the third-party services to respond, which can result in faster response times and lower error rates.

upvoted 5 times

masetromain 2 years, 5 months ago

Other options like AWS Step Functions state machine, Amazon EventBridge, and Amazon Simple Notification Service (SNS) topic are not appropriate for this use case. AWS Step Functions is a service that makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Amazon EventBridge is a serverless event bus that makes it easy to connect applications together using data from your own applications, integrated SaaS applications, and AWS services. Amazon SNS is a fully managed messaging service for both application-to-application and application-to-person (A2P) communication. These services are not focused on providing message queues and would not be the best fit for this use case.

upvoted 1 times

GabrielShiao Most Recent 5 months, 2 weeks ago

Selected Answer: A

while polling a, c is another solution accomodating the requirement. In the real case, i would pick c for a large scale eda app scenario

upvoted 1 times

AWSum1 8 months, 1 week ago

Selected Answer: A

Decoupling = SQS

upvoted 1 times

nimbus_00 8 months, 3 weeks ago

Selected Answer: A

SQS Queue = Decoupling the service calls + Eventual completion + Error handling and retries (DLQ)

upvoted 1 times

amministrazione 10 months ago

A. Use an Amazon Simple Queue Service (Amazon SQS) queue to store events and invoke the Lambda function.

upvoted 1 times

career360guru 1 year, 3 months ago

Selected Answer: A

Option A

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: A

Option A

upvoted 2 times

HC888 1 year, 7 months ago

Selected Answer: A

SQS support dead letter queue and retry if the event processed fails

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A no brainer

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: A

step functions would not help on the decoupling if you are not using an asynchronous element in this architecture which is SQS. the application need to have the ability to move out from synchronous calls to the third party services. correct answer is A.

upvoted 2 times

hpipit 2 years, 3 months ago

Selected Answer: A

A : SQS QUEUE

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: A

SQS for decoupling

upvoted 2 times

c73bf38 2 years, 4 months ago

Selected Answer: A

SQS ---> Lambda is the correct option

upvoted 2 times

zozza2023 2 years, 5 months ago

Selected Answer: A

decouple ==> SQS

upvoted 1 times

Untamables 2 years, 5 months ago

Selected Answer: A

The application needs to pass the initiative to the next step. That means the application does not wait the response from the Lambda function, it should have the responsibility only to call the Lambda function. To do so, the application only throw the job information to Amazon SQS queue and finish. After that, AWS Lambda function can pull the job information from SQS queue and start processing actively.
https://docs.aws.amazon.com/lambda/latest/dg/invocation-async.html

upvoted 2 times

Qing 2 years, 5 months ago

I vote for C - use Step Functions with its callback feature to throttle the third party api call.

upvoted 1 times

Question #101 *Topic 1*

A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.

The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sates team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.

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

A. Create a new S3 bucket in the marketing account. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing account. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket.

B. Create an SCP to grant access to the S3 bucket to the marketing account. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sates account with the marketing account. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.

C. Update the S3 bucket policy in the marketing account to grant access to the QuickSight role. Create a KMS grant for the encryption key that is used in the S3 bucket. Grant decrypt access to the QuickSight role. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.

D. Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight rote, to create a trust relationship with the new IAM role in the sales account.

Suggested Answer: *D*

Community vote distribution

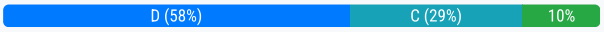

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: D

The correct answer is D. Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role to create a trust relationship with the new IAM role in the sales account.

This solution meets the requirements by allowing the marketing team to access the data in the S3 bucket in the sales account through assuming an IAM role, which eliminates the need to copy the data or share the KMS key, and also eliminates the need to modify the S3 bucket policy or create a KMS grant. This solution allows to use the same access to the bucket without duplicating data and re-encrypting it.

upvoted 27 times

masetromain 2 years, 5 months ago

A. Create a new S3 bucket in the marketing account. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing account. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket is not correct because it would create unnecessary data duplication and increased storage costs.

B. Create an SCP to grant access to the S3 bucket to the marketing account. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sales account with the marketing account. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket is not correct because it does not provide a secure way to share the KMS key between accounts and also it would create unnecessary data duplication and increased storage costs.

upvoted 4 times

masetromain 2 years, 5 months ago

C. Update the S3 bucket policy in the marketing account to grant access to the QuickSight role. Create a KMS grant for the encryption key that is used in the S3 bucket. Grant decrypt access to the QuickSight role. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket is not correct because the Sales team's S3 bucket is in a different account, so the Marketing team cannot update the policy on the Sales team's S3 bucket.

upvoted 2 times

Maria2023 Highly Voted 2 years ago

Selected Answer: D

The catch is in the answers - "Update the S3 bucket policy in the marketing account". We don't need to access a bucket in the marketing but the sales account.

upvoted 9 times

kylix75 Most Recent 5 months ago

Selected Answer: D

The correct answer is D.

Rationale:
- Lowest operational overhead using native IAM mechanisms
- Enables secure cross-account access through role assumption
- Maintains centralized access control
- No data duplication or additional storage costs
- Works seamlessly with existing KMS encryption

Other options' drawbacks:
A: Duplicates data and costs
B: SCPs aren't for granular access control
C: Incorrect bucket policy location (bucket is in sales account, not marketing)

upvoted 1 times

bhanus 6 months, 1 week ago

Selected Answer: C

Option C provides the most straightforward and efficient solution with the least operational overhead. It directly addresses the cross-account access need while maintaining security through appropriate S3 bucket and KMS key policies.

upvoted 3 times

nimbus_00 8 months, 3 weeks ago

Selected Answer: C

Creating an IAM role in the sales account that grants access to the S3 bucket and allowing the marketing account (QuickSight) to assume that role.

upvoted 1 times

amministrazione 10 months ago

C. Update the S3 bucket policy in the marketing account to grant access to the QuickSight role. Create a KMS grant for the encryption key that is used in the S3 bucket. Grant decrypt access to the QuickSight role. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.

upvoted 1 times

Jason666888 10 months, 4 weeks ago

Selected Answer: C

There must be a typo in C.

In the context of option D, if Amazon QuickSight needs to access data in an S3 bucket in a different AWS account, and the setup involves assuming multiple roles, this approach could be problematic. QuickSight would not be able to assume the role in the sales account while simultaneously using its own role in the marketing account.

upvoted 4 times

helloworldabc 10 months ago

just D

upvoted 1 times

Jason666888 10 months, 4 weeks ago

In C, "Update the S3 bucket policy in the marketing account" should be changed to "Update the S3 bucket policy in the sales account"

upvoted 3 times

8693a49 11 months ago

Selected Answer: A

What is QuickSight rote? It can't be D. I'm assuming there is no typo, so C is wrong too. B is wrong because you can't grant that permission with SCPs.

A would work provided that the replication permissions are set up correctly. It's not great because I don't think it's necessary to duplicate the data, but it's the only viable option we are given.

upvoted 1 times

Jason666888 10 months, 4 weeks ago

Dude, do you have any idea of what Petabytes amount of data mean? No one would do that in real life if there's other options
upvoted 1 times

vip2 1 year ago

Selected Answer: C

C should be correct if change typo from market account to sales account for S3 bucket policy statement.
upvoted 3 times

quizzical_kiwi 1 year ago

Selected Answer: C

Agree with other answers on C. This question is clearly a typo, and "marketing" should be changed to "sales" in C. The resolution for this scenario is even stated in the AWS Knowledge base, and the solution is identical when replacing "marketing" with "sales": https://repost.aws/knowledge-center/quicksight-cross-account-s3
upvoted 3 times

teo2157 1 year, 2 months ago

Selected Answer: C

It should be C and there should be a misspelling in "Update the S3 bucket policy in the marketing account" when it's referring to sales account
upvoted 2 times

djeong95 1 year, 3 months ago

I think this is a great question with poorly phrased answers. If I have to choose between C and D, it would be neither since they both do not provide complete answers. Let me explain:

For C, you are updating the S3 bucket policy for the marketing account, when you should be doing that for the sales account. So, C is wrong. However, if that were fixed to the sales account, everything would make sense, since the sales account would be providing the right policy, granting the correct KMS key permission, and the marketing account would be tweaking its permission in QuickSight.

For D, it is wrong simply because it says nothing about providing KMS key grant. Not only do you have to establish trust policy in the QuickSight role to access S3 bucket, you have to allow Decrypt to happen. You have to explicitly spell this out (read the permission part in the link below).

https://docs.aws.amazon.com/AmazonS3/latest/userguide/UsingKMSEncryption.html
upvoted 2 times

djeong95 1 year, 3 months ago

https://repost.aws/knowledge-center/quicksight-cross-account-s3
upvoted 1 times

VerRi 1 year, 4 months ago

Selected Answer: D

Option C: Update the S3 bucket policy in the "marketing account" ....lol
upvoted 2 times

8608f25 1 year, 4 months ago

Selected Answer: C

The answer is C. Update the S3 bucket policy in the sales account to grant access to the QuickSight role in the marketing account. Create a KMS grant for the encryption key that is used in the S3 bucket. Grant decrypt access to the QuickSight role. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.

Option C correctly identifies the need to update the S3 bucket policy to grant access specifically to the QuickSight IAM role in the marketing account, which directly addresses the requirement for cross-account access to S3 data. Additionally, creating a KMS grant for the encryption key to allow decrypt access by the QuickSight role aligns with best practices for secure, cross-account access to encrypted S3 data. This approach minimizes operational overhead by using existing roles and permissions without the need for replication or additional resource sharing mechanisms.
upvoted 2 times

AimarLeo 1 year, 4 months ago

the question is badly formulated.. with all given options missing each a spec .. none of the answers are fully convincing
upvoted 2 times

tmlong18 1 year, 5 months ago

Selected Answer: C

All answers are wrong:

A. No KMS, not necessary replication

B. No IAM

D. No KMS

But the most likely answer is C.

"Update the S3 bucket policy in the marketing account"

The question was never asked marketing s3 team bucket and all the data store in sales team S3 bucket.

I think it's a typing error (marketing-> sales).

upvoted 4 times

career360guru 1 year, 6 months ago

Selected Answer: D

Option D

upvoted 2 times

Question #102 *Topic 1*

A company is planning to migrate its business-critical applications from an on-premises data center to AWS. The company has an on-premises installation of a Microsoft SQL Server Always On cluster. The company wants to migrate to an AWS managed database service. A solutions architect must design a heterogeneous database migration on AWS.

Which solution will meet these requirements?

A. Migrate the SQL Server databases to Amazon RDS for MySQL by using backup and restore utilities.

B. Use an AWS Snowball Edge Storage Optimized device to transfer data to Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

C. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MySQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.

D. Use AWS DataSync to migrate data over the network between on-premises storage and Amazon S3. Set up Amazon RDS for MySQL. Use S3 integration with SQL Server features, such as BULK INSERT.

Suggested Answer: *C*

Community vote distribution

C (100%)

xplusfb Highly Voted 1 year, 10 months ago

Selected Answer: C

This question quietly smell weird to me but no problem answer is C
Exp : AWS Schema Conversion Tool (SCT) can automatically convert the database schema from Microsoft SQL Server to Amazon RDS for MySQL. This allows for a smooth transition of the database schema without any manual intervention. AWS DMS can then be used to migrate the data from the on-premises databases to the newly created Amazon RDS for MySQL instance. This service can perform a one-time migration of the data or can set up ongoing replication of data changes to keep the on-premises and AWS databases in sync.

upvoted 8 times

nimbus_00 Most Recent 8 months, 3 weeks ago

Selected Answer: C

Schema conversion required!

upvoted 1 times

amministrazione 10 months ago

C. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MySQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.

upvoted 1 times

TonytheTiger 1 year, 3 months ago

Selected Answer: C

This process becomes easier with services like AWS DMS and AWS Schema Conversion Tool (AWS SCT), which help you migrate your commercial database to an open-source database on AWS with minimal downtime.

In heterogeneous database migrations, the source and target databases engines are different, as in Oracle to Amazon Aurora, or Oracle to PostgreSQL, MySQL, or MariaDB migrations. The schema structure, data types, and database code in the source and target databases can be quite different, so the schema and code must be transformed before the data migration starts. For this reason, heterogeneous migration is a two-step process:

Step 1. Convert the source schema and code to match that of the target database. You can use AWS SCT for this conversion.

Step 2. Migrate data from the source database to the target database. You can use AWS DMS for this process.
https://docs.aws.amazon.com/prescriptive-guidance/latest/migration-oracle-database/heterogeneous-migration.html

upvoted 3 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: C

My 2 cents, Heterogeneous database migration and SCT go with each other

upvoted 3 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C of course

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: C

keyword = AWS Schema Conversion Tool

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: C

The question is about heterogenous database migration so in this case we need to convert the DB to a new schema.
Therefore, answer is C

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: C

Use the AWS Schema Conversion Tool

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

For heterogenous DBs, SCT is apt.

upvoted 1 times

Appon 2 years, 4 months ago

Selected Answer: C

https://aws.amazon.com/blogs/database/migrating-a-sql-server-database-to-a-mysql-compatible-database-engine/

upvoted 2 times

Musk 2 years, 4 months ago

Selected Answer: C

heterogenous -> frmo onee DB engine to another

upvoted 2 times

MasterP007 2 years, 4 months ago

Straightforward - C

upvoted 2 times

zozza2023 2 years, 5 months ago

Selected Answer: C

C is the answer

upvoted 3 times

masetromain 2 years, 5 months ago

Selected Answer: C

The correct answer is C. Use the AWS Schema Conversion Tool to translate the database schema to Amazon RDS for MySQL. Then use AWS Database Migration Service (AWS DMS) to migrate the data from on-premises databases to Amazon RDS.

AWS Schema Conversion Tool (SCT) can automatically convert the database schema from Microsoft SQL Server to Amazon RDS for MySQL. This allows for a smooth transition of the database schema without any manual intervention.

AWS DMS can then be used to migrate the data from the on-premises databases to the newly created Amazon RDS for MySQL instance. This service can perform a one-time migration of the data or can set up ongoing replication of data changes to keep the on-premises and AWS databases in sync.

upvoted 4 times

masetromain 2 years, 5 months ago

Option A is not correct because while Amazon RDS for MySQL supports SQL Server databases, it is not a good fit for migrating business-critical applications. The data model and architecture are different and would require significant re-engineering.

Option B is not correct because AWS Snowball Edge Storage Optimized devices are used for transferring large amounts of data to and from AWS, but they do not support SQL Server.

Option D is not correct because AWS DataSync can only transfer files and folders, it does not support SQL Server databases.

upvoted 2 times

Question #103 *Topic 1*

A publishing company's design team updates the icons and other static assets that an ecommerce web application uses. The company serves the icons and assets from an Amazon S3 bucket that is hosted in the company's production account. The company also uses a development account that members of the design team can access.

After the design team tests the static assets in the development account, the design team needs to load the assets into the S3 bucket in the production account. A solutions architect must provide the design team with access to the production account without exposing other parts of the web application to the risk of unwanted changes.

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

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket.

B. In the development account, create a new IAM policy that allows read and write access to the S3 bucket.

C. In the production account, create a role Attach the new policy to the role. Define the development account as a trusted entity.

D. In the development account, create a role. Attach the new policy to the role Define the production account as a trusted entity.

E. In the development account, create a group that contains all the IAM users of the design team Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role In the production account.

F. In the development account, create a group that contains all the IAM users of the design team Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the development account.

Suggested Answer: *ADE*

Community vote distribution

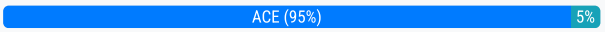

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: ACE

The correct answer is A, C, and E.

A: In the production account, creating a new IAM policy that allows read and write access to the S3 bucket is correct because it allows the design team to upload and update the static assets in the S3 bucket in the production account.

C: In the production account, creating a role and attaching the new policy to the role, and defining the development account as a trusted entity is correct because it allows the design team from the development account to assume the role and access the S3 bucket in the production account, while limiting their access to only the specific resources and actions defined in the policy.

upvoted 14 times

masetromain 2 years, 5 months ago

E: In the development account, creating a group that contains all the IAM users of the design team and attaching a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account is correct because it allows the users in the group to assume the role created in the production account, which gives them access to the S3 bucket in the production account.

The other choices are not correct because:

B: In the development account, creating a new IAM policy that allows read and write access to the S3 bucket is not correct because the design team needs to access the S3 bucket in the production account, not the development account.

upvoted 4 times

masetromain 2 years, 5 months ago

D: In the development account, creating a role, attaching the new policy to the role and defining the production account as a trusted entity is not correct because the design team needs to assume a role in the production account to access the S3 bucket, not create a role in the development account.

F: In the development account, creating a group that contains all the IAM users of the design team and attaching a different IAM policy to the

group to allow the sts:AssumeRole action on the role in the development account is not correct because the design team needs to assume a role in the production account to access the S3 bucket, not the development account.

upvoted 2 times

zejou1 Highly Voted 2 years, 3 months ago

Selected Answer: ACE

Step 1: Create a role in the Production Account; create the role in the Production account and specify the Development account as a trusted entity. You also limit the role permissions to only read and write access to the productionapp bucket. Anyone granted permission to use the role can read and write to the productionapp bucket.
Step 2: Grant access to the role Sign in as an administrator in the Development account and allow the AssumeRole action on the UpdateApp role in the Production account.

So, recap, production account you create the policy for S3, and you set development account as a trusted entity. Then on the development account you allow the sts:assumeRole action on the role in production account.
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

upvoted 10 times

LuongTo 7 months ago

C: "creating a role and attaching the new policy to the role" => it is very clear to use the policy to control read write. A question about the role created with C, where to use?

upvoted 1 times

amministrazione Most Recent 10 months ago

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket.
D. In the development account, create a role. Attach the new policy to the role Define the production account as a trusted entity.
E. In the development account, create a group that contains all the IAM users of the design team Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role In the production account.

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: ACE

ACE. F is a trap.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: ACE

A, C and E

upvoted 1 times

AMohanty 1 year, 7 months ago

BCE
Need to provide Account in Dev S3 Read Write Access
We define the permissions of the user in the Account it was created in

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: ACE

ACE in this case

upvoted 1 times

MoussaNoussa 2 years ago

ACE is the correct choice of course

upvoted 1 times

leehjworking 2 years, 1 month ago

Selected Answer: ACE

Vote for ACE

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: ACE

ACE is the best choice

upvoted 3 times

God_Is_Love 2 years, 3 months ago

Selected Answer: ACE

Make Dev account as trusted entity. create a role in prod account. attache IAM policy of prod account and let development account assume this role to access prod s3 bucket.

upvoted 2 times

Musk 2 years, 4 months ago

Selected Answer: ACE

I think it's clear

upvoted 1 times

tatdatpham 2 years, 4 months ago

Selected Answer: ACE

ACE is correct answer

upvoted 2 times

zozza2023 2 years, 5 months ago

Selected Answer: ACE

ACE should works

upvoted 2 times

zhangyu20000 2 years, 5 months ago

ACE is my answer

upvoted 2 times

masetromain 2 years, 5 months ago

Selected Answer: ADE

A, D, and E are the correct steps that would meet the requirements.

A. In the production account, create a new IAM policy that allows read and write access to the S3 bucket. This will allow the design team to read and write to the S3 bucket that holds the assets in the production account.

D. In the development account, create a role. Attach the new policy to the role. Define the production account as a trusted entity. This will allow the design team to assume a role in the development account that has permissions to access the S3 bucket in the production account.

E. In the development account, create a group that contains all the IAM users of the design team. Attach a different IAM policy to the group to allow the sts:AssumeRole action on the role in the production account. This will allow the users in the design team group to assume the role created in step D and access the S3 bucket in the production account.

upvoted 2 times

masetromain 2 years, 5 months ago

Option B is not required because the design team needs to access the S3 bucket in the production account, not in the development account.

Option C is not required because the design team needs to access the S3 bucket in the production account and this can be done by assuming a role in the development account.

Option F is not required because the design team needs to access the S3 bucket in the production account and this can be done by assuming a role in the development account that is trusted by the production account.

upvoted 1 times

Question #104 *Topic 1*

A company developed a pilot application by using AWS Elastic Beanstalk and Java. To save costs during development, the company's development team deployed the application into a single-instance environment. Recent tests indicate that the application consumes more CPU than expected. CPU utilization is regularly greater than 85%, which causes some performance bottlenecks.

A solutions architect must mitigate the performance issues before the company launches the application to production.

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

A. Create a new Elastic Beanstalk application. Select a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the maximum CPU utilization is over 85% for 5 minutes.

B. Create a second Elastic Beanstalk environment. Apply the traffic-splitting deployment policy. Specify a percentage of incoming traffic to direct to the new environment in the average CPU utilization is over 85% for 5 minutes.

C. Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.

D. Select the Rebuild environment action with the load balancing option. Select an Availability Zones. Add a scale-out rule that will run if the sum CPU utilization is over 85% for 5 minutes.

Suggested Answer: *A*

Community vote distribution

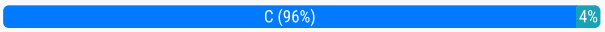

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: C

I think AWS wants you to know is the below.
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

upvoted 27 times

ninomfr64 Highly Voted 1 year, 5 months ago

A = you don't need to create a new application (instead you could create a new environment in the existing application)
B = traffic-split is used to deploy a new version of the app, not to scale out
C = correct
D = rebuild does not allow to change environment configuration https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environment-management-rebuild.html

upvoted 5 times

nimbus_00 Most Recent 8 months, 3 weeks ago

Selected Answer: C

You can change your environment type to a single-instance or load-balanced, scalable environment by editing your environment's configuration.

upvoted 2 times

amministrazione 10 months ago

C. Modify the existing environment's capacity configuration to use a load-balanced environment type. Select all Availability Zones. Add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes.

upvoted 1 times

Maygam 1 year, 6 months ago

Selected Answer: C

You can change the existing environment from single instance to load balanced.
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

upvoted 3 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C

upvoted 1 times

yuliaqwerty 1 year, 6 months ago

C here https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/GettingStarted.EditConfig.html

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: C

you can change existing Beanstalk environment type from a single instance to load-balanced

upvoted 2 times

CuteRunRun 1 year, 10 months ago

Selected Answer: C

I prefer C

upvoted 1 times

Spaco 1 year, 11 months ago

Selected Answer: C

Option C is very correct. See https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html for confirmation

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

its a C

upvoted 1 times

leehjworking 2 years, 1 month ago

Anybody know why we should select all AZs?

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: C

Modify the existing environment's capacity configuration to use a load-balanced environment type.

upvoted 1 times

zejou1 2 years, 3 months ago

Selected Answer: C

You can change your environment type to a single-instance or load-balanced, scalable environment by editing your environment's configuration. In some cases, you might want to change your environment type from one type to another. For example, let's say that you developed and tested an application in a single-instance environment to save costs. When your application is ready for production, you can change the environment type to a load-balanced, scalable environment so that it can scale to meet the demands of your customers.
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

upvoted 4 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

A is wrong. no need to re create new EB env when the question is asking to mitigate probable performance issues based on current compute consumption of >=85%

upvoted 2 times

spd 2 years, 4 months ago

Selected Answer: C

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

upvoted 2 times

Musk 2 years, 4 months ago

Selected Answer: C

It's C. https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html#using-features.managing.changetype

upvoted 1 times

Question #105 *Topic 1*

A finance company is running its business-critical application on current-generation Linux EC2 instances. The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.

Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?

A. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.

B. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month.

C. Using Amazon CloudWatch with AWS Lambda to change the type, size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric.

D. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

Suggested Answer: *B*

Community vote distribution

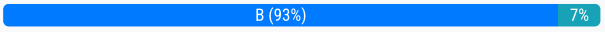

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

B. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month.

This is the optimal solution as migrating the database to Amazon RDS will provide the ability to easily scale read replicas for handling increased read traffic during the end of the month. Additionally, RDS will manage the underlying infrastructure and provide automatic backups, software patching, and monitoring, which will reduce the operational overhead for the company.

Option A may help but it will not be sufficient to handle the heavy load, option C and D are not efficient solutions to han

upvoted 16 times

amministrazione Most Recent 10 months ago

B. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: B

B, include read replicas

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: B

Option B -> Reporting workload = higher read operation ==> Solution RDS read replica.

upvoted 1 times

hansean 1 year, 8 months ago

Selected Answer: D

I go with D

upvoted 1 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: D

I vote D

To solve heavy IO issue, I think both option B and D both works.But the question demands for to "handle the month-end load with the LEAST impact

on performance" , Option B create the new read replicas during end of month seems too complicated, you'll need to seperate read/write traffic from application at the end of the month.

upvoted 1 times

venvig 1 year, 10 months ago

Selected Answer: B

Reporting is also an important hint. Only read operations are needed here; so read replicas would server the purpose

upvoted 2 times

xplusfb 1 year, 10 months ago

Selected Answer: B

all other sections not applicable i guess specially D its so funny. Each month none of technical person doesnt want to do like this task.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

B of cpourse

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

it slows down during the final three days of each month due to month-end reporting

then

hight read in database == solution add read replicas

B

upvoted 2 times

nexus2020 1 year, 11 months ago

month end reporting is to submit the financial data, aka write the new data to DB

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: B

Performing a one-time migration

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: B

B is the best solution

upvoted 2 times

Question #106 *Topic 1*

A company runs a Java application that has complex dependencies on VMs that are in the company's data center. The application is stable. but the company wants to modernize the technology stack. The company wants to migrate the application to AWS and minimize the administrative overhead to maintain the servers.

Which solution will meet these requirements with the LEAST code changes?

A. Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Container. Store container images in Amazon Elastic Container Registry (Amazon ECR). Grant the ECS task execution role permission 10 access the ECR image repository. Configure Amazon ECS to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

B. Migrate the application code to a container that runs in AWS Lambda. Build an Amazon API Gateway REST API with Lambda integration. Use API Gateway to interact with the application.

C. Migrate the application to Amazon Elastic Kubernetes Service (Amazon EKS) on EKS managed node groups by using AWS App2Container. Store container images in Amazon Elastic Container Registry (Amazon ECR). Give the EKS nodes permission to access the ECR image repository. Use Amazon API Gateway to interact with the application.

D. Migrate the application code to a container that runs in AWS Lambda. Configure Lambda to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

Suggested Answer: *B*

Community vote distribution

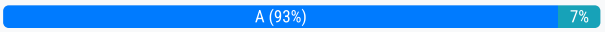

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

The correct answer would be A, as migrating the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Container and storing container images in Amazon Elastic Container Registry (Amazon ECR) would minimize the code changes and administrative overhead required to maintain the servers. This option would allow the company to use the Application Load Balancer (ALB) to interact with the application and the ECS task execution role permission to access the ECR image repository.

Option B would require the application code to be migrated to a container that runs in AWS Lambda, which would require more code changes.

Option C would require migrating the application to Amazon Elastic Kubernetes Service (Amazon EKS) which would require more administrative overhead.

Option D would require configuring Lambda to use an Application Load Balancer (ALB), which is not a native feature of Lambda.

upvoted 20 times

Musk 2 years, 4 months ago

B does not say anything about Lambda. Where have you red that?

upvoted 1 times

Musk 2 years, 4 months ago

You are right, I mixed A with B

upvoted 1 times

rbm2023 2 years, 1 month ago

There is another problem with Option B, it suggest using EKS with managed node groups and not Fargate, which breaks the requirement for reducing administrative overhead

upvoted 1 times

masetromain 2 years, 5 months ago

This solution allows for the existing application code to be packaged into a container, which can then be deployed to ECS on Fargate. The use of AWS App2Container will help automate the containerization process, minimizing the need for code changes. Additionally, by using ECR to store container images, the application can continue to use the same images and dependencies that it currently relies on. The use of an Application Load Balancer (ALB) to interact with the application further simplifies the migration process by allowing the use of the existing application's endpoint.

upvoted 4 times

zejou1 Highly Voted 2 years, 3 months ago

Selected Answer: A

AWS App2Container (A2C) is a command line tool to help you lift and shift applications that run in your on-premises data centers or on virtual machines, so that they run in containers that are managed by Amazon ECS, Amazon EKS, or AWS App Runner.

Moving legacy applications to containers is often the starting point toward application modernization. There are many benefits to containerization:
• Reduces operational overhead and infrastructure costs
• Increases development and deployment agility
• Standardizes build and deployment processes across an organization
https://docs.aws.amazon.com/app2container/latest/UserGuide/what-is-a2c.html
AWS Fargate is a serverless, pay-as-you-go compute engine that lets you focus on building applications without managing servers. AWS Fargate is compatible with both Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS).
https://aws.amazon.com/fargate/

upvoted 9 times

amministrazione Most Recent 10 months ago

A. Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Container. Store container images in Amazon Elastic Container Registry (Amazon ECR). Grant the ECS task execution role permission 10 access the ECR image repository. Configure Amazon ECS to use an Application Load Balancer (ALB). Use the ALB to interact with the application.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A, ECS Fargate

upvoted 1 times

AimarLeo 1 year, 5 months ago

Selected Answer: A

If the keyword 'Java' has not been mentioned, Answer A would have been considered as A2C (App2Container) is valid only for Java and .Net web applications

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: A

A = correct
B = migrating app to container to be executed in a Lambda requires more code changes
C = EKS with managed node group requires more operations than ECS with Fargate
D = see B

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: A

Option A. Option C EKS not not valid because as using API Gateway is not needed and may require more code changes.

upvoted 2 times

severlight 1 year, 7 months ago

Selected Answer: A

in the case of fargate capacity provider you should grant permissions to access ecr to task execution role, otherwise to ec2 instance roles which you run containers on

upvoted 1 times

CVDON 1 year, 8 months ago

Sorry is A

upvoted 1 times

CVDON 1 year, 8 months ago

C on eks because of complex VM dependecies

upvoted 1 times

CVDON 1 year, 8 months ago

D because of complex vm dependencies

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

it's an A

upvoted 1 times

Maria2023 2 years ago

Did anyone notice that part "has complex dependencies on VMs that are in the company's data center."? If the application has complex dependencies on VMs then how do we migrate it to containers or lambda? Another awkward question.

upvoted 1 times

Sarutobi 2 years, 2 months ago

Selected Answer: A

I still select A, but as someone that has migrated Java applications to AWS using AWS App2Container and RedHat S2i, this is a lot of pain.

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: A

Migrate the application to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate by using AWS App2Container.

upvoted 1 times

kiran15789 2 years, 3 months ago

Selected Answer: A

least code chansges

upvoted 2 times

keonlee 2 years, 4 months ago

Selected Answer: A

Fargate, Modernize stack

upvoted 2 times

Question #107 *Topic 1*

A company has an asynchronous HTTP application that is hosted as an AWS Lambda function. A public Amazon API Gateway endpoint invokes the Lambda function. The Lambda function and the API Gateway endpoint reside in the us-east-1 Region. A solutions architect needs to redesign the application to support failover to another AWS Region.

Which solution will meet these requirements?

A. Create an API Gateway endpoint in the us-west-2 Region to direct traffic to the Lambda function in us-east-1. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

B. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure API Gateway to direct traffic to the SQS queue instead of to the Lambda function. Configure the Lambda function to pull messages from the queue for processing.

C. Deploy the Lambda function to the us-west-2 Region. Create an API Gateway endpoint in us-west-2 10 direct traffic to the Lambda function in us-west-2. Configure AWS Global Accelerator and an Application Load Balancer to manage traffic across the two API Gateway endpoints.

D. Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

Suggested Answer: *B*

Community vote distribution

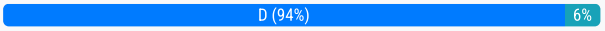

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: D

The correct answer is D. Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints. This solution meets the requirement of having a failover to another region by having a copy of the Lambda function and API Gateway endpoint in a different region, and using Route 53's failover routing policy to route traffic between the two regions.

Option A is not correct because it only creates an additional API Gateway endpoint in us-west-2 and relies on Route 53's failover routing policy to direct traffic to the correct endpoint. But it does not deploy the Lambda function to the new region and this makes the failover incomplete.

upvoted 23 times

testingaws123 2 years, 3 months ago

You always use ChatGPT to paste answers. Most of the time ChatGPT gives wrong answers do you know this?

upvoted 12 times

juanife 4 months, 2 weeks ago

If the answer explanation of why it's one option and why the other ones are not ok truly represents the correct answer then I would not say anything. I think chat gpt is very useful if you (with knowledge on mind) are able to judge what this ai machine says and validate that.

upvoted 1 times

masetromain 2 years, 5 months ago

Option B is not correct because it uses a SQS queue as a buffer between the API Gateway and the Lambda function, but this does not provide failover to another region. In addition, it would also increase the latency of the system as the SQS will act as an additional layer.

Option C is not correct because it deploys the Lambda function to the us-west-2 Region and creates an API Gateway endpoint in the same region. But it uses AWS Global Accelerator and an Application Load Balancer to manage traffic across the two API Gateway endpoints. However, this is not a failover solution as both regions will be active and serving traffic at the same time.

upvoted 3 times

amministrazione Most Recent 10 months ago

D. Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region. Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: D

D, deploy everything in the second region and configure the failover routing policy

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: D

Option D

upvoted 1 times

venvig 1 year, 10 months ago

Selected Answer: D

Refer https://aws.amazon.com/blogs/architecture/implementing-multi-region-disaster-recovery-using-event-driven-architecture/

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

clearly D

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: D

Deploy the Lambda function and an API Gateway endpoint to the us-west-2 Region

upvoted 1 times

zejou1 2 years, 3 months ago

Selected Answer: D

Currently, the default API endpoint type in API Gateway is the edge-optimized API endpoint, which enables clients to access an API through an Amazon CloudFront distribution. This typically improves connection time for geographically diverse clients. By default, a custom domain name is globally unique and the edge-optimized API endpoint would invoke a Lambda function in a single region in the case of Lambda integration. You can't use this type of endpoint with a Route 53 active-active setup and fail-over.

The new regional API endpoint in API Gateway moves the API endpoint into the region and the custom domain name is unique per region. This makes it possible to run a full copy of an API in each region and then use Route 53 to use an active-active setup and failover.
https://aws.amazon.com/blogs/compute/building-a-multi-region-serverless-application-with-amazon-api-gateway-and-aws-lambda/

upvoted 2 times

God_Is_Love 2 years, 3 months ago

Selected Answer: D

B is wrong, cannot direct traffic to SQS Queue ? it does not even mention posting messages to queue.

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: D

The correct answer is D

upvoted 2 times

zhangyu20000 2 years, 5 months ago

D is correct

A is not because the Lambda is in us-ease-1 but api gateway is in us-west-2. cannot cross regions

upvoted 4 times

masetromain 2 years, 5 months ago

Selected Answer: A

The correct answer is A.

In this solution, an API Gateway endpoint is created in the us-west-2 Region. This new endpoint is configured to direct traffic to the Lambda function in us-east-1. If a failure occurs in the us-east-1 Region, Amazon Route 53's failover routing policy automatically routes traffic to the us-west-2 Region. This ensures that traffic is directed to a healthy endpoint, providing failover support for the application.

B, C and D does not meet the requirement of having failover routing policy.

In B, SQS is not a failover mechanism, it is a messaging service and it does not provide failover routing.

In C, Global Accelerator and Application Load Balancer does not provide failover routing.

In D, While creating a second endpoint in the us-west-2 Region and using Amazon Route 53 to route traffic to it, it still does not provide failover routing.
upvoted 2 times

CProgrammer 1 year, 9 months ago
D CLEARLY States: Configure Amazon Route 53 to use a failover routing policy to route traffic for the two API Gateway endpoints. You claimed it did not , and the moderator ALLOWED IT ?!? !?
upvoted 1 times

CProgrammer 1 year, 9 months ago
Gateway VPC endpoints provide reliable connectivity to Amazon S3 and DynamoDB without requiring an internet gateway or a NAT device for your VPC.

https://docs.aws.amazon.com/vpc/latest/privatelink/gateway-endpoints.html

==> IN CONTRAST
These are the ENDPOINTS for API Gateway:

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-endpoint-types.html
Gateway endpoint DOES NOT DIRECT TRAFFIC PERIOD
upvoted 1 times

Question #108 *Topic 1*

A retail company has structured its AWS accounts to be part of an organization in AWS Organizations. The company has set up consolidated billing and has mapped its departments to the following OUs: Finance, Sales, Human Resources (HR), Marketing, and Operations. Each OU has multiple AWS accounts, one for each environment within a department. These environments are development, test, pre-production, and production.

The HR department is releasing a new system that will launch in 3 months. In preparation, the HR department has purchased several Reserved Instances (RIs) in its production AWS account. The HR department will install the new application on this account. The HR department wants to make sure that other departments cannot share the RI discounts.

Which solution will meet these requirements?

A. In the AWS Billing and Cost Management console for the HR department's production account turn off RI sharing.

B. Remove the HR department's production AWS account from the organization. Add the account 10 the consolidating billing configuration only.

C. In the AWS Billing and Cost Management console. use the organization's management account 10 turn off RI Sharing for the HR departments production AWS account.

D. Create an SCP in the organization to restrict access to the RIs. Apply the SCP to the OUs of the other departments.

Suggested Answer: *C*

Community vote distribution

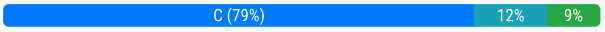

kiran15789 Highly Voted 2 years, 3 months ago

Selected Answer: C

Management account -->Billing Dashboard --> Billing preferences, this option is there to choose enable/disable RI discounts sharing

upvoted 9 times

amministrazione Most Recent 10 months ago

C. In the AWS Billing and Cost Management console. use the organization's management account 10 turn off RI Sharing for the HR departments production AWS account.

upvoted 2 times

Dgix 1 year, 3 months ago

Selected Answer: C

It is indeed C.

upvoted 2 times

Dgix 1 year, 3 months ago

Selected Answer: A

RI sharing is done for the whole Org. It's all or nothing, and it's done in the Billing and Cost Management console in the Org account.

upvoted 2 times

helloworldabc 10 months ago

just C

upvoted 1 times

JOKERO 1 year, 3 months ago

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html

C

upvoted 1 times

8608f25 1 year, 4 months ago

Selected Answer: A

Option A is correct because AWS allows the management of RI sharing settings at the account level within the AWS Billing and Cost Management console. By turning off RI sharing in the HR department's production account, the RI benefits (such as the discounted rate) are applied only to instances within that account, preventing other accounts, even within the same organization, from accessing these discounts. This directly addresses

the requirement.

Option C suggests using the organization's management account to turn off RI sharing for the HR department's production AWS account. While the management account controls many aspects of AWS Organizations, including consolidated billing, RI sharing preferences are managed at the individual account level within the Billing and Cost Management console, not directly through the management account for specific accounts.
upvoted 2 times

horyoryo 1 year, 6 months ago

Selected Answer: C

option C
upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

surely C
upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: C

C is the way to go
upvoted 1 times

sambb 2 years, 3 months ago

Selected Answer: C

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/ri-turn-off.html
upvoted 3 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

Management account -->Billing Dashboard --> Billing preferences, this option is there to choose enable/disable RI discounts sharing
https://us-east-1.console.aws.amazon.com/billing/home#/preferences
upvoted 3 times

testingaws123 2 years, 3 months ago

Selected Answer: D

How can you restrict access from AWS billing console? Can you show me please??
Option D is the correct solution because an SCP (Service Control Policy) can be created in the AWS Organizations service to restrict access to specific resources or actions across the entire organization or specific OUs. In this case, an SCP can be created to restrict other departments from accessing the RIs purchased by the HR department's production account. This ensures that the discounts are not shared with other departments.
upvoted 3 times

- **God_Is_Love** 2 years, 3 months ago
  Bro, Go to Management account -->Billing Dashboard --> Billing preferences, this option is there to choose enable/disable RI discounts sharing
  https://us-east-1.console.aws.amazon.com/billing/home#/preferences
  upvoted 5 times

- **chikorita** 1 year, 10 months ago
  initially i thought the same....but the catch here is that RIs are purchased in HR Prod department
  So, we have to work on disabling discount sharing wrt that account so that IT IS NOT SHARED W OTHERS and this actions can only be performed from Management account
  upvoted 1 times

- **SK_Tyagi** 1 year, 10 months ago
  Restricting the access to RI's is not the ask in the question, only "restricting the RI discounts" from HR to other departments is the ask, and that you could be done by Management Account (as identified by others in this forum). Hope that helps!
  upvoted 2 times

jojom19980 2 years, 4 months ago

Selected Answer: C

The correct answer is C

upvoted 1 times

masetromain 2 years, 5 months ago

Selected Answer: C

The correct answer is C.

In this solution, the organization's management account can be used to turn off RI sharing for the HR department's production AWS account in the AWS Billing and Cost Management console. This will ensure that the other departments cannot share the RI discounts and the HR department can use the RIs for their new system without any interruption.

upvoted 3 times

masetromain 2 years, 5 months ago

A, B and D does not meet the requirement of turning off RI sharing for the HR department's production AWS account.

In A, Turning off RI sharing in the HR department's production account will not prevent other departments from sharing the RI discounts.

In B, Removing the HR department's production AWS account from the organization may cause issues in consolidated billing and it does not prevent other departments from sharing the RI discounts.

In D, Creating an SCP in the organization to restrict access to the RIs is not necessary because the management account can directly turn off the RI sharing, it also does not prevent other departments from sharing the RI discounts.

upvoted 3 times

Question #109 *Topic 1*

A large company is running a popular web application. The application runs on several Amazon EC2 Linux instances in an Auto Scaling group in a private subnet. An Application Load Balancer is targeting the instances in the Auto Scaling group in the private subnet. AWS Systems Manager Session Manager is configured, and AWS Systems Manager Agent is running on all the EC2 instances.

The company recently released a new version of the application. Some EC2 instances are now being marked as unhealthy and are being terminated. As a result, the application is running at reduced capacity. A solutions architect tries to determine the root cause by analyzing Amazon CloudWatch logs that are collected from the application, but the logs are inconclusive.

How should the solutions architect gain access to an EC2 instance to troubleshoot the issue?

A. Suspend the Auto Scaling group's HealthCheck scaling process. Use Session Manager to log in to an instance that is marked as unhealthy.

B. Enable EC2 instance termination protection. Use Session Manager to log in to an instance that is marked as unhealthy.

C. Set the termination policy to OldestInstance on the Auto Scaling group. Use Session Manager to log in to an instance that is marked an unhealthy.

D. Suspend the Auto Scaling group's Terminate process. Use Session Manager to log in to an instance that is marked as unhealthy.

Suggested Answer: *D*

Community vote distribution

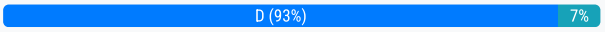

zozza2023 Highly Voted 2 years, 5 months ago

Selected Answer: D

The correct answer is D.

upvoted 10 times

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: D

Disabling health check wont let SA know which instance is un healthy. So A is certainly wrong. D is correct.

upvoted 9 times

amministrazione Most Recent 10 months ago

D. Suspend the Auto Scaling group's Terminate process. Use Session Manager to log in to an instance that is marked as unhealthy.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: D

D, stop the autoscaling process

upvoted 1 times

AWSLord32 1 year, 5 months ago

Why not B? Can the ASG override the Ec2 termination protection?

upvoted 3 times

career360guru 1 year, 6 months ago

Selected Answer: D

Option D

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: D

you can stop auto-scaling processes, here you need to stop termination, you need health checks to know which instance to check

upvoted 2 times

venvig 1 year, 10 months ago

Selected Answer: D

If ASG terminates the instances because they are unhealthy there is no way we can login to the instance using session manager or otherwise to investigate the problem. So, suspend the termination.

upvoted 4 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

d of course

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: D

keyword == Auto Scaling group's Terminate process.

upvoted 1 times

Alando 1 year, 9 months ago

Have you cleared the exam?

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: D

Suspend the Auto Scaling group's Terminate process.

upvoted 2 times

zejou1 2 years, 3 months ago

Selected Answer: D

Amazon EC2 Auto Scaling stops marking instances unhealthy as a result of EC2 and Elastic Load Balancing health checks. Your custom health checks continue to function properly. After you suspend HealthCheck, if you need to, you can manually set the health state of instances in your group and have ReplaceUnhealthy replace them.

Suspending the Terminate process doesn't prevent the successful termination of instances using the force delete option with the delete-auto-scaling-group command.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-suspend-resume-processes.html

https://docs.aws.amazon.com/systems-manager/latest/userguide/incident-manager.html

We want the health checks to continue failing, just stop terminating to identify root cause

upvoted 4 times

testingaws123 2 years, 3 months ago

Selected Answer: A

Answer is A

If you do not want instances to be replaced, we recommend that you suspend the ReplaceUnhealthy and HealthCheck process for individual Auto Scaling groups. For more information, see Suspend and resume a process for an Auto Scaling group.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/ec2-auto-scaling-health-checks.html

upvoted 3 times

zejou1 2 years, 3 months ago

That does not solve, it removes the healthcheck process, but also removes the ones that are being marked as unhealthy. The issue now is that one it is tagged as unhealthy they are being terminated. So, any that are already marked get terminated and you just removed the health checks to find remaining. you can't troubleshoot what you don't know.

upvoted 5 times

masetromain 2 years, 5 months ago

Selected Answer: D

https://www.examtopics.com/discussions/amazon/view/51249-exam-aws-certified-solutions-architect-professional-topic-1/

The correct answer is D.

In this solution, the architect can suspend the Auto Scaling group's Terminate process, which will prevent the instances marked as unhealthy from being terminated. This will allow the architect to log in to the instance using Session Manager and troubleshoot the issue without losing access to the instance.

upvoted 7 times

masetromain 2 years, 5 months ago

Option A is incorrect because suspending the HealthCheck scaling process will not prevent instances from being terminated.

Option B is incorrect because enabling EC2 instance termination protection will not prevent instances from being terminated by Auto Scaling group.

Option C is incorrect because setting the termination policy to OldestInstance on the Auto Scaling group will not prevent instances marked as unhealthy from being terminated.

upvoted 4 times

Question #110 *Topic 1*

A company wants to deploy an AWS WAF solution to manage AWS WAF rules across multiple AWS accounts. The accounts are managed under different OUs in AWS Organizations.

Administrators must be able to add or remove accounts or OUs from managed AWS WAF rule sets as needed. Administrators also must have the ability to automatically update and remediate noncompliant AWS WAF rules in all accounts.

Which solution meets these requirements with the LEAST amount of operational overhead?

A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage. Update the parameter as needed to add or remove accounts or OUs. Use an Amazon EventBridge rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account.

B. Deploy an organization-wide AWS Config rule that requires all resources in the selected OUs to associate the AWS WAF rules. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.

C. Create AWS WAF rules in the management account of the organization. Use AWS Lambda environment variables to store account numbers and OUs to manage. Update environment variables as needed to add or remove accounts or OUs. Create cross-account IAM roles in member accounts. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts.

D. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization. Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage. Update AWS KMS as needed to add or remove accounts or OUs. Create IAM users in member accounts. Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts.

Suggested Answer: *D*

Community vote distribution

A (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

The correct answer is A.

In this solution, AWS Firewall Manager is used to manage AWS WAF rules across accounts in the organization. An AWS Systems Manager Parameter Store parameter is used to store account numbers and OUs to manage. This parameter can be updated as needed to add or remove accounts or OUs. An Amazon EventBridge rule is used to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account. This solution allows for easy management of AWS WAF rules across multiple accounts with minimal operational overhead.

upvoted 20 times

masetromain 2 years, 5 months ago

Option B does not meet the requirement of being able to add or remove accounts or OUs from managed AWS WAF rule sets as needed.

Option C is not the best approach as it requires manual configuration of the cross-account IAM roles and assume-role calls in the Lambda function, increasing the operational overhead.

Option D does not meet the requirement of providing a centralized management console to manage the WAF rules across multiple accounts.

upvoted 3 times

Aquaman 3 months, 2 weeks ago

B doesn't allow you to target just accounts. The question is asking for a solution that can target accounts and OUs

upvoted 1 times

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: A

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

upvoted 6 times

d0ug7979 Most Recent 9 months, 1 week ago

Selected Answer: B

Correct answer is B. I would have said A like everyone else, but correct answer was provided in Udemy practice exam.

Thanks to Organization structure, Config rules apply automatically to newly added accounts (fulfills requirements: least amount of operational overhead (as opposed to A - manually maintaining accounts and OU list).

As often, AWS exam answers are partially off-track, a real-life deployment would be a clever combination of both A & B answers, using FW manager, Config and Cloudformation.
https://aws.amazon.com/blogs/security/use-aws-firewall-manager-to-deploy-protection-at-scale-in-aws-organizations/

upvoted 2 times

amministrazione 10 months ago

A. Use AWS Firewall Manager to manage AWS WAF rules across accounts in the organization. Use an AWS Systems Manager Parameter Store parameter to store account numbers and OUs to manage. Update the parameter as needed to add or remove accounts or OUs. Use an Amazon EventBridge rule to identify any changes to the parameter and to invoke an AWS Lambda function to update the security policy in the Firewall Manager administrative account.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: A

Option A

upvoted 1 times

venvig 1 year, 10 months ago

Selected Answer: A

AWS Firewall Manager is a security management service which allows you to centrally configure and manage firewall rules across your accounts and applications in AWS Organizations
Firewall Manager supports wide variety of services, including:

- AWS WAF
- VPC Security Groups
- AWS Network Firewall
- Route53 DNS Firewall
- AWS Shield Advanced
- Palo Alto Cloud Next-generation firewalls

The Prerequisites are: AWS Organizations + AWS Config.

upvoted 5 times

CuteRunRun 1 year, 10 months ago

Selected Answer: A

I have to say A is right.
please take a look at this:
https://aws.amazon.com/blogs/security/centrally-manage-aws-waf-api-v2-and-aws-managed-rules-at-scale-with-firewall-manager/

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A is a good option

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: A

keyword == AWS Firewall Manager

upvoted 3 times

tromyunpak 2 years ago

the correct answer is A https://docs.aws.amazon.com/solutions/latest/automations-for-aws-firewall-manager/architecture-overview.html

upvoted 2 times

rbm2023 2 years, 1 month ago

Selected Answer: A

This is a complex question. But I voted A because the Firewall manager seems to be the correct way to centralize the rules across accounts. Below are some interesting references I could find
https://catalog.us-east-1.prod.workshops.aws/workshops/4cbaea3b-ceba-48e3-bd56-eca138f7a66c/en-US
https://aws.amazon.com/blogs/security/use-aws-firewall-manager-vpc-security-groups-to-protect-applications-hosted-on-ec2-instances/

https://aws.amazon.com/blogs/security/automatically-updating-aws-waf-rule-in-real-time-using-amazon-eventbridge/

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: A

Use AWS Firewall Manager to manage AWS WAF rules

upvoted 2 times

God_Is_Love 2 years, 3 months ago

Selected Answer: A

Not D, KMS to store account numbers ?

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: A

The correct answer is A.

upvoted 2 times

Question #111 *Topic 1*

A solutions architect is auditing the security setup or an AWS Lambda function for a company. The Lambda function retrieves, the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.

The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the Internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.

What should the solutions architect recommend to meet these requirements?

A. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

B. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Enforce HTTPS on the connection to Amazon S3 during data transfers.

C. Save the database credentials in AWS Systems Manager Parameter Store. Set up password rotation on the credentials in Parameter Store. Change the IAM role for the Lambda function to allow the function to access Parameter Store. Modify the Lambda function to retrieve the credentials from Parameter Store. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

D. Save the database credentials in AWS Secrets Manager. Set up password rotation on the credentials in Secrets Manager. Change the IAM role for the Lambda function to allow the function to access Secrets Manager. Modify the Lambda function to retrieve the credentials from Secrets Manager. Enforce HTTPS on the connection to Amazon S3 during data transfers.

Suggested Answer: *D*

Community vote distribution

zozza2023 Highly Voted 2 years, 5 months ago

Selected Answer: A

a little bit confused between A and D but as said by others members D doesn't adress the The question of "data must not travel across the Internet"==> A is the answer

upvoted 20 times

MikelH93 Highly Voted 2 years, 2 months ago

Selected Answer: A

B and D are out because you need the VPC endpoints.

C is out because you cannot enable rotation in Parameter Store

(https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_parameterstore.html)

upvoted 7 times

Paul123456789 Most Recent 2 months, 4 weeks ago

Selected Answer: D

Option A does not address the requirement for rotating database credentials.

upvoted 1 times

amministrazione 10 months ago

A. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

upvoted 1 times

MAZIADI 10 months, 2 weeks ago

Selected Answer: A

A is better than D because it remove the complexity of management of the secret to connect to the DB and replaces it with the IAM DB authentication. In addition S3 endpoint GW is better to prevent traffic going through internet.

upvoted 1 times

AWSPro1234 1 year, 3 months ago

Selected Answer: A

Key is data must not travel accros the internet mean use VPC gateway
upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A, "data must no travel across the internet". This setup ensures internal network use only, meeting the security and networking requirements efficiently
upvoted 1 times

a54b16f 1 year, 4 months ago

Selected Answer: A

The data must not travel across the Internet.
upvoted 2 times

8608f25 1 year, 4 months ago

Selected Answer: D

Option D offers a comprehensive solution by leveraging AWS Secrets Manager for storing and automatically rotating database credentials, which directly addresses the concern of minimizing the impact if credentials become compromised. Changing the Lambda function to retrieve credentials from Secrets Manager enhances security by not storing credentials within environment variables. Enforcing HTTPS for S3 data transfers ensures the data in transit is encrypted. While deploying a gateway VPC endpoint for S3 (as mentioned in other options) is a best practice to keep traffic within the AWS network, enforcing HTTPS also contributes to securing data transfers without explicitly stating the need to avoid Internet travel. Secrets Manager inherently provides secure access to secrets without needing to travel across the Internet when accessed from AWS services within the same region.

Option A does not address the requirement for securing and rotating database credentials stored as Lambda environment variables.
upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: A

Answer is A as S3 VPC is endpoint is needed to avoid data going over internet.
upvoted 1 times

task_7 1 year, 9 months ago

Selected Answer: D

AWS Secrets Manager is meant for this job, why go with any other option
upvoted 2 times

task_7 1 year, 9 months ago

My bad its A
D is not addressing this point
The data must not travel across the Internet
upvoted 6 times

CuteRunRun 1 year, 10 months ago

I prefor A
upvoted 2 times

Jonalb 1 year, 11 months ago

Selected Answer: A

A

https://aws.amazon.com/pt/blogs/database/iam-role-based-authentication-to-amazon-aurora-from-serverless-applications/
upvoted 3 times

Jonalb 1 year, 12 months ago

Selected Answer: D

https://docs.aws.amazon.com/pt_br/secretsmanager/latest/userguide/vpc-endpoint-overview.html
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A for sure
upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: A

I was about to chose D however just enforcing the HTTP will not avoid the data to travel across internet. You will need the option where the gateway VPC endpoint is deployed for access the S3. The answer is A.

A will also solve the issue related to authenticate the lambda to aurora without needing to store passwords, refer to - https://aws.amazon.com/blogs/database/iam-role-based-authentication-to-amazon-aurora-from-serverless-applications/

upvoted 1 times

OCHT 2 years, 2 months ago

Selected Answer: D

However, Option A is not the best choice for the given scenario because:

It doesn't address the requirement to minimize the impact of compromised database credentials. IAM database authentication eliminates traditional user credentials, but it doesn't implement password rotation for the remaining IAM credentials.

While the VPC endpoint keeps traffic within the AWS network, it doesn't enforce encryption during data transfers to Amazon S3.

Option D, on the other hand, addresses both the requirement of minimizing the impact of compromised credentials through password rotation using AWS Secrets Manager and ensuring encrypted data transfers to Amazon S3 by enforcing HTTPS. That's why Option D is the better choice for this scenario.

upvoted 3 times

rbm2023 2 years, 1 month ago

I was also choosing D however just enforcing the HTTP will not avoid the data to travel across internet. You will need the option where the gateway VPC endpoint is deployed for access the S3. The answer is A

upvoted 3 times

Question #112 *Topic 1*

A large mobile gaming company has successfully migrated all of its on-premises infrastructure to the AWS Cloud. A solutions architect is reviewing the environment to ensure that it was built according to the design and that it is running in alignment with the Well-Architected Framework.

While reviewing previous monthly costs in Cost Explorer, the solutions architect notices that the creation and subsequent termination of several large instance types account for a high proportion of the costs. The solutions architect finds out that the company's developers are launching new Amazon EC2 instances as part of their testing and that the developers are not using the appropriate instance types.

The solutions architect must implement a control mechanism to limit the instance types that only the developers can launch.

Which solution will meet these requirements?

A. Create a desired-instance-type managed rule in AWS Config. Configure the rule with the instance types that are allowed. Attach the rule to an event to run each time a new EC2 instance is launched.

B. In the EC2 console, create a launch template that specifies the instance types that are allowed. Assign the launch template to the developers' IAM accounts.

C. Create a new IAM policy. Specify the instance types that are allowed. Attach the policy to an IAM group that contains the IAM accounts for the developers

D. Use EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden image.

Suggested Answer: *C*

Community vote distribution

C (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct answer is C.

In this solution, a new IAM policy is created that specifies the allowed instance types. This policy is then attached to an IAM group that contains the IAM accounts for the developers. This will ensure that the developers can only launch instances of the specified types, thus limiting the costs associated with the creation and termination of large instances.

upvoted 15 times

masetromain 2 years, 5 months ago

A. Creating a desired-instance-type managed rule in AWS Config is not a sufficient solution, as it only identifies when an instance is launched with an unauthorized type, it does not prevent it.

B. Creating a launch template that specifies the instance types that are allowed is not a sufficient solution, because it limits the instances types that can be launched in the EC2 console, but it does not prevent the launch of instances through the AWS SDK, AWS CLI, or other AWS services.

D. Using EC2 Image Builder to create an image pipeline for the developers and assist them in the creation of a golden image is not a direct solution to the problem of limiting the instance types that only the developers can launch. It can be useful for creating standardize images for the developers, but it does not provide the necessary control mechanism to limit the instance types.

upvoted 12 times

gagol14 Highly Voted 1 year, 5 months ago

Selected Answer: C

{
"Sid": "limitedSize",
"Effect": "Deny",
"Action": "ec2:RunInstances",
"Resource": "arn:aws:ec2:*:*:instance/*",
"Condition": {
"ForAnyValue:StringNotLike": {

"ec2:InstanceType": [
"*.nano",
"*.small",
"*.micro",
"*.medium"
]
}
}
}
upvoted 6 times

amministrazione Most Recent 10 months ago
C. Create a new IAM policy. Specify the instance types that are allowed. Attach the policy to an IAM group that contains the IAM accounts for the developers
upvoted 1 times

cox1960 1 year, 5 months ago
"an IAM group that contains the IAM accounts" ???
upvoted 1 times

igor12ghsj577 1 year, 5 months ago
yes, in IAM group you have user IAM accounts.
upvoted 1 times

sse69 1 year ago
You have IAM users...Not IAM "accounts". Bad wording here...
upvoted 2 times

career360guru 1 year, 6 months ago
Selected Answer: C
Option C
upvoted 1 times

NikkyDicky 1 year, 12 months ago
Selected Answer: C
Its a C
upvoted 1 times

Maria2023 2 years ago
Selected Answer: C
The only technical achievable choices are A and C. However A will only identify the issue and will not prevent it. Even if we set up a remediation rule to terminate the instances immediately - that will cause more issues for the developers and unclear signals that something is wrong with the testing. So A remains the only possible option.
upvoted 2 times

Parimal1983 2 years ago
C is the correct solution remained. Typo mistake in the comments.
upvoted 1 times

easytoo 2 years ago
c-c-c-c-c-cc-c-c-cc-c-c-c-c-cc-
upvoted 1 times

mfsec 2 years, 3 months ago
Selected Answer: C
IAM policy..
upvoted 1 times

zozza2023 2 years, 5 months ago
Selected Answer: C
answer is C
upvoted 3 times

Question #113 *Topic 1*

A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost for cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.

Which actions should a solutions architect lake to resolve the problem and prevent it from happening in the future? (Choose three.)

A. Create an AWS Config rule in each account to find resources with missing tags.

B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.

C. Use Amazon Inspector in the organization to find resources with missing tags.

D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.

E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.

F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.

Suggested Answer: *CDE*

Community vote distribution

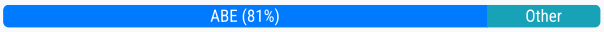

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: ABE

If config rule is added (A) it can be seen in AWS Config aggregator (E) Using SCP in as aws organization is used here in question. So, A,B,E

upvoted 7 times

God_Is_Love 2 years, 3 months ago

If there are no organizations used, D can be used to prevent EC2 run instances too,

C is for vulnerabilities checking..F for all security issues consolidated..

upvoted 4 times

OCHT Highly Voted 2 years, 2 months ago

Selected Answer: ABE

A. Create an AWS Config rule in each account to find resources with missing tags.

By creating an AWS Config rule in each account, you can check if resources are missing tags or have tags that are not conforming to your organization's standards. You can also use AWS Config to automatically remediate non-compliant resources by applying tags. This can help ensure that resources are properly tagged for cost allocation purposes. Here is the AWS Config documentation for creating rules:

https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_use-managed-rules.html

upvoted 5 times

OCHT 2 years, 2 months ago

E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.

By creating an AWS Config aggregator, you can collect a list of EC2 instances across multiple accounts in the organization that are missing the required Project tag. This can help you identify instances that need to be tagged properly for cost allocation. Here is the AWS Config documentation for creating aggregators:

https://docs.aws.amazon.com/config/latest/developerguide/config-aggregator.html

upvoted 7 times

AWSLord32 1 year, 5 months ago

So what is the point of having A if you have E at an Org level?

upvoted 2 times

fartosh 1 year, 1 month ago

AWS Config aggregator does not run any rules on its own. Instead, it collects the data from the "source accounts" where AWS Config is enabled.

A to get the list of EC2 instances in each account.

E to aggregate the lists from all accounts in one place.

B to disallow creating non-compliant EC2 instances.

See https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html.

upvoted 3 times

OCHT 2 years, 2 months ago

B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.

By creating a Service Control Policy (SCP) in the organization, you can enforce a deny action for EC2 instances that do not have the required Project tag. This can prevent users from launching instances that are not tagged correctly and ensure that new instances are tagged properly for cost allocation. Here is the AWS Organizations documentation for creating SCPs:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html

upvoted 5 times

amministrazione Most Recent 10 months ago

A. Create an AWS Config rule in each account to find resources with missing tags.

B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.

E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: ABE

ABE, SCP + Config + Config Aggregator

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: BE

B and E handle the requirements in a centralised manner, giving least operational overhead, without anything needing to be added. The question is plainly wrongly stated. If three options have to be selected, then A is the least absurd one.

upvoted 3 times

8608f25 1 year, 4 months ago

Selected Answer: ABE

A. Create an AWS Config rule in each account to find resources with missing tags.AWS Config can evaluate the configuration of your AWS resources and identify resources that do not comply with specified requirements, such as missing specific tags. This helps in identifying existing resources with the issue.

B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.Service Control Policies (SCPs) can enforce permissions across all accounts in an organization. By creating an SCP that denies launching EC2 instances without the required Project tag, you can prevent the problem from occurring in the future at the organization level.

E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.An AWS Config aggregator can aggregate compliance data from multiple accounts and regions. This allows for centralized visibility of instances lacking the required tags, making it easier to address and resolve the issue across the entire organization.

upvoted 2 times

AWSLord32 1 year, 5 months ago

Selected Answer: BDE

A is not needed if you have D. Correct answer is BDE.

upvoted 1 times

AWSLord32 1 year, 5 months ago

I meant E, not D

upvoted 1 times

8608f25 1 year, 4 months ago

It is not D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.

IAM policies do not directly support conditional denies based on tag presence during the resource creation process in the same way SCPs do. This enforcement is better handled at the organization level with SCPs.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: ABE

Option A, B and E

upvoted 1 times

Sandeep_B 1 year, 8 months ago

Selected Answer: ABE

Inspector checks for Vulnerabilities but not the tags.

upvoted 3 times

NikkyDicky 1 year, 12 months ago

Selected Answer: ABE

its ABE

upvoted 2 times

youngmanaws 2 years, 2 months ago

A. 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. These documents define the actions to be performed on noncompliant AWS resources evaluated by AWS Config Rules. You can associate SSM documents by using AWS Management Console or by using APIs.

AWS Config provides a set of managed automation documents with remediation actions. You can also create and associate custom automation documents with AWS Config rules.

To apply remediation on noncompliant resources, you can either choose the remediation action you want to associate from a prepopulated list or create your own custom remediation actions using SSM documents. AWS Config provides a recommended list of remediation action in the AWS Management Console.

In the AWS Management Console, you can either choose to manually or automatically remediate noncompliant resources by associating remediation actions with AWS Config rules. With all remediation actions, you can either choose manual or automatic remediation.

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: ABE

ABE is the better choice

upvoted 1 times

Damijo 2 years, 3 months ago

what's the value of A and E together- it's either or ? the outcome is the same - thoughts?

upvoted 4 times

AWSLord32 1 year, 5 months ago

Fully agree, BDE

upvoted 1 times

AWSLord32 1 year, 5 months ago

Did some research..

Aggregators provide a read-only view into the source accounts and regions that the aggregator is authorized to view. Aggregators do not provide mutating access into the source account or region. For example, this means that you cannot deploy rules through an aggregator or pull snapshot files from the source account or region through an aggregator.
https://docs.aws.amazon.com/config/latest/developerguide/config-concepts.html#multi-account-multi-region-data-aggregation

So ABE seems correct

upvoted 2 times

jaysparky 2 years, 4 months ago

ABE makes sense

upvoted 1 times

spd 2 years, 4 months ago

Selected Answer: ABE

Config, SCP and IAM policy may not require in each account but it says to select three options so going with ABE

upvoted 1 times

Musk 2 years, 4 months ago

Selected Answer: AE

BE makes sense

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: ABE

the best way to deploy config rules accross accounts= SCP

upvoted 2 times

Question #114 *Topic 1*

A company has an on-premises monitoring solution using a PostgreSQL database for persistence of events. The database is unable to scale due to heavy ingestion and it frequently runs out of storage.

The company wants to create a hybrid solution and has already set up a VPN connection between its network and AWS. The solution should include the following attributes:
• Managed AWS services to minimize operational complexity.
• A buffer that automatically scales to match the throughput of data and requires no ongoing administration.
• A visualization tool to create dashboards to observe events in near-real time.
• Support for semi-structured JSON data and dynamic schemas.

Which combination of components will enable the company to create a monitoring solution that will satisfy these requirements? (Choose two.)

A. Use Amazon Kinesis Data Firehose to buffer events. Create an AWS Lambda function to process and transform events.

B. Create an Amazon Kinesis data stream to buffer events. Create an AWS Lambda function to process and transform events.

C. Configure an Amazon Aurora PostgreSQL DB cluster to receive events. Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards.

D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events. Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards.

E. Configure an Amazon Neptune DB instance to receive events. Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards.

Suggested Answer: *AD*

Community vote distribution

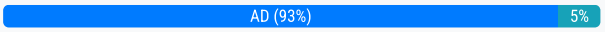

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: AD

Amazon Kinesis Data Firehose (A) allows you to buffer events in two ways: through buffering size or buffering time. With buffering size, you can configure the maximum size of the buffer in MB or the maximum number of records in the buffer. Once the buffer is full, it will automatically deliver the data to the destination

Amazon ES (D) has its ability to receive events from various sources in real-time. Amazon ES can ingest data from a variety of sources, such as Amazon Kinesis Data Firehose, Amazon CloudWatch Logs, and Amazon S3, making it a powerful tool for organizations looking to analyze and visualize real-time streaming data. (Kibana dashboards)

upvoted 14 times

OCHT Highly Voted 2 years, 2 months ago

Selected Answer: AD

Option B includes using an Amazon Kinesis data stream to buffer events, which is a valid solution for a streaming data use case. However, it requires more ongoing administration compared to using Amazon Kinesis Data Firehose, which is a fully managed service. Additionally, the use of Amazon Kinesis Data Firehose allows the company to take advantage of built-in data transformation and processing capabilities, which can reduce the amount of code required to implement the solution. Therefore, I selected option A over option B as it better meets the requirement of minimizing operational complexity.

upvoted 12 times

Paul123456789 Most Recent 2 months, 4 weeks ago

Selected Answer: BD

AWS Lambda is a source for Amazon Kinesis Data Firehose not a destination
https://docs.aws.amazon.com/firehose/latest/dev/create-name.html
https://docs.aws.amazon.com/firehose/latest/dev/create-destination.html
also, Firehose encountered timeout errors when calling AWS Lambda. The maximum supported function timeout is 5 minutes
https://docs.aws.amazon.com/firehose/latest/dev/data-transformation.html
correct answer B and D

upvoted 2 times

albert_kuo 3 months, 4 weeks ago

Selected Answer: BD

Option B (Kinesis Data Streams + Lambda) + Option D (Amazon ES + Kibana):
Buffer: Kinesis Data Streams automatically scales and buffers events.
Processing: Lambda transforms JSON events and sends them to Amazon ES.
Storage: Amazon ES stores semi-structured JSON with dynamic schemas.
Visualization: Kibana provides near-real-time dashboards.
Managed: All services (Kinesis, Lambda, ES, Kibana) are fully managed.
Workflow: Events flow from on-premises via VPN to Kinesis → Lambda → ES → Kibana.
Result: Meets all requirements seamlessly.

upvoted 2 times

GabrielShiao 5 months, 2 weeks ago

Selected Answer: BD

While most voted AD, I vote BD. I picked B instead of A because you can not use lambda to access the kinesis firehose directly.

upvoted 2 times

albert_kuo 3 months, 4 weeks ago

Agreeed

upvoted 1 times

amministrazione 10 months ago

A. Use Amazon Kinesis Data Firehose to buffer events. Create an AWS Lambda function to process and transform events.
D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events. Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards.

upvoted 1 times

Smart 1 year, 1 month ago

Selected Answer: AD

"A buffer that automatically scales to match the throughput of data and requires no ongoing administration."
I think buffer, here, means a solution that will reliably hold information for further successful processing. I don't think it means to buffer and batch process the events so I don't agree with other people's comments in regards to buffer.

That said, my concern is with "automatically scales to match the throughput of data". Firehose does it automatically. Kinesis can also do automatically if on-demand mode is chosen.

Also, "Support for semi-structured JSON data and dynamic schemas." Dynamic Schemas? Firehose or Data stream don't do that. Firehose does do dynamic partitioning and JSON deserializing. I guess that's what the question meant?

upvoted 1 times

TonytheTiger 1 year, 2 months ago

Selected Answer: AD

Option A NOT Option B - Amazon Data Firehose buffers incoming streaming data in memory to a certain size (buffering size) and for a certain period of time (buffering interval) before delivering it to the specified destinations.

https://docs.aws.amazon.com/firehose/latest/dev/buffering-hints.html

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: AD

On second thought: A because B requires manual shard configuration.

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: BE

Also, Streams is more real-time.

upvoted 1 times

AWSum1 8 months, 4 weeks ago

It says "Near Real Time" not "Real Time" so Firehouse is the better option between the 2

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: BD

B rather than A because B integrates the lambda functionality for transformation of the data, which must be done as an added step in A, thereby increasing operational overhead.

upvoted 2 times

8608f25 1 year, 4 months ago

Selected Answer: AD

A. Use Amazon Kinesis Data Firehose to buffer events. Create an AWS Lambda function to process and transform events.Amazon Kinesis Data Firehose provides a fully managed service for effortlessly loading streaming data into AWS services such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk. It scales automatically to match the throughput of data and requires no ongoing administration. AWS Lambda can be used in conjunction with Kinesis Data Firehose to process and transform the data before it's loaded into the destination, supporting dynamic schemas and semi-structured JSON data. Additionally, Amazon Kinesis Data Firehose has built-in buffering capabilities and can be used to observe events in near-real time, making it a more appropriate choice for the given scenario.

upvoted 2 times

8608f25 1 year, 4 months ago

D. Configure Amazon Elasticsearch Service (Amazon ES) to receive events. Use the Kibana endpoint deployed with Amazon ES to create near-real-time visualizations and dashboards. Amazon Elasticsearch Service (Amazon ES) is a managed service that makes it easy to deploy, secure, operate, and scale Elasticsearch to search, analyze, and visualize data in real-time. Kibana is an open-source visualization tool designed to work with Elasticsearch, providing powerful and easy-to-use features to create dashboards that can visualize data in near-real-time.

upvoted 1 times

AimarLeo 1 year, 5 months ago

ElasticSearch is the ex name of new OpenSearch

upvoted 2 times

ninomfr64 1 year, 5 months ago

Selected Answer: BD

I choose Data Stream (KDS) over Data Firehose (KDF) in this scenario:

- KDS allows to you store events up to 1 year, allowing to achieve buffering with no constraints on size and with a very large time limit. KDS support on-demand capacity mode
- KDF transport mechanism is based on buffering, but here buffering is limited on size (max 128MiB) and time (up to 900 sec)

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: AD

A and D

upvoted 1 times

AMohanty 1 year, 9 months ago

BD

Question states near-Real time

Thats the differentiating factor between Kinesis data stream and Firehose

I would go for B and D

upvoted 2 times

chikorita 1 year, 9 months ago

but about "• Managed AWS services to minimize operational complexity."

i believe Kinesis Firehose is managed solution whereas DataStream required operational overhead

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: AD

AD for unstructured data

upvoted 1 times

Question #115 *Topic 1*

A team collects and routes behavioral data for an entire company. The company runs a Multi-AZ VPC environment with public subnets, private subnets, and in internet gateway. Each public subnet also contains a NAT gateway. Most of the company's applications read from and write to Amazon Kinesis Data Streams. Most of the workloads run in private subnets.

A solutions architect must review the infrastructure. The solution architect needs to reduce costs and maintain the function of the applications. The solutions architect uses Cost Explorer and notices that the cost in the EC2-Other category is consistently high. A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

What should the solutions architect do to meet these requirements?

A. Enable VPC Flow Logs. Use Amazon Athena to analyze the logs for traffic that can be removed. Ensure that security groups are blocking traffic that is responsible for high costs.

B. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that applications have the correct IAM permissions to use the interface VPC endpoint.

C. Enable VPC Flow Logs and Amazon Detective. Review Detective findings for traffic that is not related to Kinesis Data Streams. Configure security groups to block that traffic.

D. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that the VPC endpoint policy allows traffic from the applications.

Suggested Answer: *D*

Community vote distribution

D (93%) 8%

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: D

VPC endpoints to mitigate NAT gateway huge data transfer costs especially in Kinesis usecase where large data is passed thru

With a VPC endpoint policy, you can define rules to control access to the VPC endpoint. You can specify the source IP address or IP address range that is allowed to access the endpoint, as well as the type of traffic that is allowed, such as HTTP, HTTPS, or custom TCP ports. You can also specify the resources that can be accessed through the VPC endpoint, such as an Amazon S3 bucket or an Amazon DynamoDB table.

upvoted 14 times

Maria2023 Highly Voted 2 years ago

Selected Answer: D

B is a distractor. You don't need IAM permissions to use a service via an endpoint. You only need to set up proper routing to that endpoint

upvoted 10 times

kylix75 Most Recent 5 months ago

Selected Answer: B

The correct answer is B.

Rationale:

Interface VPC endpoints for Kinesis eliminate NAT gateway traffic costs

Applications in private subnets can access Kinesis through the VPC endpoint

IAM permissions are the proper security control

Maintains functionality while reducing costs

Other options issues:

A/C: Flow logs analysis won't reduce NAT costs

D: Similar to B but focuses on endpoint policy instead of IAM permissions

upvoted 2 times

Heman31in 6 months, 2 weeks ago

Selected Answer: D

Why Option D is a Better Fit Here
Cost Reduction Goal: The scenario is primarily about reducing NAT gateway costs by using a VPC endpoint. A properly configured VPC endpoint policy ensures applications can connect to Kinesis through the private endpoint without hitting NAT gateways.
IAM Permissions Likely Already Exist: If the applications are already interacting with Kinesis, their IAM permissions should already be in place. The focus, therefore, shifts to configuring the new VPC endpoint properly.
Endpoint Policy Completeness: VPC endpoint policies act as a resource-based policy at the network level, which is critical for ensuring that applications can route their traffic correctly through the VPC endpoint.
upvoted 1 times

youonebe 7 months, 1 week ago
Answer is B.

Option D is incorrect.
While similar to B, focuses on endpoint policy instead of IAM permissions
VPC endpoint policies alone are insufficient
IAM permissions are crucial for application access
upvoted 1 times

Syre 9 months, 1 week ago
Selected Answer: B
Access Permissions are still required for most AWS services, including Kinesis Data Streams, even when accessed via a VPC endpoint. The endpoint allows traffic to the service, but your application or users still need IAM permissions to interact with the service. Without proper IAM permissions, even if the routing is set up correctly, the service will not authorize actions like reading from or writing to a Kinesis stream.
upvoted 2 times

amministrazione 10 months ago
D. Add an interface VPC endpoint for Kinesis Data Streams to the VPC. Ensure that the VPC endpoint policy allows traffic from the applications.
upvoted 1 times

red_panda 1 year, 2 months ago
Selected Answer: D
D without any doubt.
upvoted 1 times

gofavad926 1 year, 3 months ago
Selected Answer: D
D, VPC endpoint
upvoted 2 times

gofavad926 1 year, 3 months ago
Selected Answer: D
D, VPC endpoint
upvoted 1 times

career360guru 1 year, 6 months ago
Selected Answer: D
Option D
upvoted 1 times

rlf 1 year, 8 months ago
Answer is D.
An endpoint policy is a resource-based policy that you attach to a VPC endpoint to control which AWS principals can use the endpoint to access an AWS service.
https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-access.html
upvoted 1 times

NikkyDicky 1 year, 11 months ago
Selected Answer: D
It's a d
upvoted 1 times

SkyZeroZx 2 years ago
Selected Answer: D
reduce cost == interface VPC endpoint

upvoted 3 times

SkyZeroZx 2 years ago

A further review shows that NatGateway-Bytes charges are increasing the cost in the EC2-Other category.

upvoted 1 times

Anonymous9999 2 years, 2 months ago

Selected Answer: D

D is the answer.

It's not B because user's/applications doesn't need permissions to use an endpoint:
https://docs.aws.amazon.com/vpc/latest/privatelink/security_iam_id-based-policy-examples.html

upvoted 2 times

romiao106 2 years, 1 month ago

No. in your document it says "By default, users and roles don't have permission to create or modify AWS PrivateLink resources". Users and roles don't have permissions so they do need permissions to use an interface endpoint

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: D

D is the best choice.

upvoted 1 times

Sarutobi 2 years, 3 months ago

If this is a cost-saving question is very hard to answer, you pay for both, and depending on the region one can be cheaper than the other. There is a cost for a NAT GW and also for a VPCendpoint per AZ plus the traffic you generate over them. In my experience, because you need a VPCendpoint for each service NAT-GW is cheaper.

upvoted 1 times

fartosh 1 year, 1 month ago

I agree that both NAT GW and interface VPC endpoints can become expensive. I believe that's why the question mentioned that most applications use KDS. I assume that it's the biggest middleware service and you will not need VPC endpoints for other services.

Pricing (based on Ohio):
NAT GW: 0.045 $/h + 0.045 $/GB
Interface VPC Endpoint: 0.01 $/h + 0.01 $/GB (lowered if more data transferred)

In the final setup the company will still pay for NAT GW (hourly fee) but the transfer cost (most of it) will be moved to VPCE, which gives:
for 1 GB per month
NAT GW: (24*30)h*0.045$/h + 1GB*0.045$/GB = 32.445$ > (24*30)h*0.01$/h + 1GB*0.01$/GB = 7.21$
for 1000GB per month
NAT GW: (24*30)h*0.045$/h + 1000GB*0.045$/GB = 77.4$ > (24*30)h*0.01$/h + 1000GB*0.01$/GB = 17.2$

upvoted 1 times

Question #116 *Topic 1*

A retail company has an on-premises data center in Europe. The company also has a multi-Region AWS presence that includes the eu-west-1 and us-east-1 Regions. The company wants to be able to route network traffic from its on-premises infrastructure into VPCs in either of those Regions. The company also needs to support traffic that is routed directly between VPCs in those Regions. No single points of failure can exist on the network.

The company already has created two 1 Gbps AWS Direct Connect connections from its on-premises data center. Each connection goes into a separate Direct Connect location in Europe for high availability. These two locations are named DX-A and DX-B, respectively. Each Region has a single AWS Transit Gateway that is configured to route all inter-VPC traffic within that Region.

Which solution will meet these requirements?

A. Create a private VIF from the DX-A connection into a Direct Connect gateway. Create a private VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with the Direct Connect gateway. Peer the transit gateways with each other to support cross-Region routing.

B. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Associate the eu-west-1 transit gateway with this Direct Connect gateway. Create a transit VIF from the DX-8 connection into a separate Direct Connect gateway. Associate the us-east-1 transit gateway with this separate Direct Connect gateway. Peer the Direct Connect gateways with each other to support high availability and cross-Region routing.

C. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway. Configure the Direct Connect gateway to route traffic between the transit gateways.

D. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway. Peer the transit gateways with each other to support cross-Region routing.

Suggested Answer: *A*

Community vote distribution

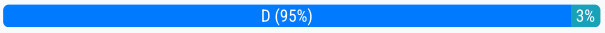

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: D

https://docs.aws.amazon.com/images/whitepapers/latest/hybrid-connectivity/images/dx-dxgw-transit-gateway-multi-region-public-vif.png
B is wrong as it says, two DX Gateways contradictory
C is wrong as it says to configure DXG to route traffic. infact Transit gateway peering need to be done between two transit gateways of each reigon.
A is wrong because Private VIF is not apt in mentioned config of the question. Public VIF is correct (Transit public VIF)
If you are using a single DX Gateway

upvoted 16 times

God_Is_Love 2 years, 3 months ago

Whichever option has this text is correct - "Peer the transit gateways with each other to support cross-Region routing"

upvoted 5 times

Syre Most Recent 9 months, 1 week ago

Selected Answer: C

While transit gateway peering can enable cross-Region VPC communication, it is not necessary when you are using a Direct Connect gateway. A Direct Connect gateway already provides the capability to route traffic across multiple Regions without needing to peer the transit gateways directly.

upvoted 2 times

amministrazione 10 months ago

D. Create a transit VIF from the DX-A connection into a Direct Connect gateway. Create a transit VIF from the DX-B connection into the same Direct Connect gateway for high availability. Associate both the eu-west-1 and us-east-1 transit gateways with this Direct Connect gateway. Peer the transit gateways with each other to support cross-Region routing.

upvoted 1 times

teo2157 1 year, 2 months ago

It can be both A or D based on AWS documentation: https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/hybrid-network-connections.html

upvoted 2 times

Dgix 1 year, 3 months ago

Selected Answer: D

Don't let "No single points of failure can exist on the network" mislead you into thinking that you need two DCGWs. DCGWs are not part of the region they connect to. Therefore, no SPOF translates to a double DC connection to a single DCGW. Hence, D.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: D

D, this approach ensures high availability and robust network connectivity across the specified AWS regions and the on-premises data center.

upvoted 1 times

Jassybanga 1 year, 4 months ago

Answer D - As per from AWS
https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/aws-dx-dxgw-with-aws-transit-gateway-multi-regions-more-than-3.html

upvoted 3 times

career360guru 1 year, 6 months ago

Selected Answer: D

Choice is between C and D. Better the two D is the right option.

upvoted 1 times

subbupro 1 year, 6 months ago

D is correct ref architecture https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/aws-dx-dxgw-with-aws-transit-gateway-multi-regions-and-aws-public-peering.html

upvoted 4 times

mnsait 7 months ago

This is the best answer I found for this question. Thank you @subbupro for the reference. It explains exactly what is needed to understand here.

upvoted 1 times

shaaam80 1 year, 6 months ago

Selected Answer: D

Answer D. Peer the transit gateways for cross-region routing.

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: D

to connect to transit gateways through the dx gateway you should use transit VIF

upvoted 1 times

frfavoreto 1 year, 9 months ago

I agree 'D' is a good answer to the problem, but isn't the DXGW a single point of failure?

Question says "No single points of failure can exist on the network."

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

it's D

upvoted 1 times

happystrawberry 2 years, 1 month ago

Would it be C for the answer? A Direct Connect gateway supports communication between attached transit virtual interfaces and associated transit gateways only and may enable a virtual private gateway to another virtual private gateway.
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-transit-gateways.html

upvoted 1 times

happystrawberry 2 years, 1 month ago

Actually, D is a proper answer.

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: D

I agree with option D

Refer to the diagram below which explains in detail the use of Transit VIF and Public VIF. Also demonstrates the necessity for peering the transit gateways to allow the cross-region routing.

https://docs.aws.amazon.com/images/whitepapers/latest/hybrid-connectivity/images/dx-dxgw-transit-gateway-multi-region-public-vif.png

The only options that are using the cross-region routing are A and D. Option A mentions the use of Private VIF and not the Transit VIF. Hence A is incorrect.

upvoted 4 times

rbm2023 2 years, 1 month ago

Refer to the following article

https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/aws-dx-dxgw-with-aws-transit-gateway-multi-regions-and-aws-public-peering.html

upvoted 3 times

dev112233xx 2 years, 2 months ago

Selected Answer: D

Transit VIF required to connect to Transit Gateway, and Transit peering is required to connect multi regions...

Here is the full diagram:

https://docs.aws.amazon.com/whitepapers/latest/hybrid-connectivity/aws-dx-dxgw-with-aws-transit-gateway-multi-regions-and-aws-public-peering.html

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: D

D is the answer

upvoted 2 times

Question #117 *Topic 1*

A company is running an application in the AWS Cloud. The company's security team must approve the creation of all new IAM users. When a new IAM user is created, all access for the user must be removed automatically. The security team must then receive a notification to approve the user. The company has a multi-Region AWS CloudTrail trail in the AWS account.

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

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.

B. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.

C. Invoke a container that runs in Amazon Elastic Container Service (Amazon ECS) with AWS Fargate technology to remove access.

D. Invoke an AWS Step Functions state machine to remove access.

E. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.

F. Use Amazon Pinpoint to notify the security team.

Suggested Answer: *ADE*

Community vote distribution

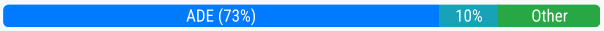

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: ADE

Event Bus (EventBridge) system to receive event notification (Option A). Step function can get triggered with workflow of doing steps like removing access and sending email etc..(Option D, E)

EventBridge enables you to create event rules that match events from different sources, such as AWS services, SaaS applications, custom applications, and other AWS accounts. Once an event rule is triggered, EventBridge can route the event to one or more targets, such as AWS Lambda functions, Amazon SNS topics, Amazon SQS queues, or custom HTTP endpoints.

AWS Step Functions supports several AWS services, such as AWS Lambda, Amazon Simple Notification Service (SNS), and Amazon Simple Queue Service (SQS). You can use these services to trigger actions and pass data between steps in your state machine.

Pinpoint is chat system which question did not ask, F is wrong. Not C as

upvoted 14 times

Jay_2pt0_1 2 years, 1 month ago

I agree with this.

upvoted 1 times

hobokabobo 2 years, 3 months ago

this explanation makes sense to me.

upvoted 1 times

sergza888 Most Recent 1 week, 5 days ago

Selected Answer: ACE

I Just do not see any reasoning to use step functions. There is no orchestration is involved here. you Just need lambda or Fargate to remove the access. D does not say anything about lambda step at all i think it is just distractor

upvoted 1 times

amministrazione 10 months ago

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.

D. Invoke an AWS Step Functions state machine to remove access.

E. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.

upvoted 1 times

TonytheTiger 1 year, 2 months ago

Selected Answer: ACE

Option ADE: Most people agree with option AE. There can be situations where human intervention is required before the workflow can progress. For example, approving a substantial credit increase may require human approval

https://docs.aws.amazon.com/step-functions/latest/dg/use-cases-security-automation.html

upvoted 1 times

helloworldabc 10 months ago

just ADE

upvoted 1 times

24Gel 1 year, 3 months ago

Why not BCE? or ACE?

How to use Step Function to remove permission?

upvoted 1 times

dankositzke 1 year, 4 months ago

Poorly constructed answer choices, but ADE is the least worst option.

upvoted 2 times

zanhsieh 1 year, 4 months ago

Selected Answer: ADE

I picked ADE. EventBridge, Lambda / Step Function, and SNS are required.

BDE: No. CloudTrail can't trigger Step Function directly.

ABE: No. This solution can't remove the user access automatically.

Choosing B alone without A can't directly trigger Lambda / Step functions to remove the user access. C can't compare with D. F is not relevant.

upvoted 1 times

AWSLord32 1 year, 5 months ago

Selected Answer: BDE

Eventbridge is not needed. Cloudtrail can send notifications to SNS directly

https://docs.aws.amazon.com/awscloudtrail/latest/userguide/configure-sns-notifications-for-cloudtrail.html

upvoted 3 times

altonh 5 months, 2 weeks ago

"You can be notified when CloudTrail publishes new log files to your Amazon S3 bucket. You manage notifications using Amazon Simple Notification Service (Amazon SNS)."

That's the only notification you are getting. It's not good enough. You need the actual API call made, which is the user creation API.

upvoted 1 times

AWSLord32 1 year, 5 months ago

Also, if you select ADE how would the event ever trigger SNS to send the notification?

upvoted 2 times

fartosh 1 year, 1 month ago

What do you mean? SNS topic is one of the (many) allowed targets for EventBridge.

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

Regarding "Eventbridge is not needed" - it's only true for notifications because CloudTrail integrates with SNS. CloudTrail alone cannot trigger any automation tools like Lambda or Step Function. That's why EventBridge is better in this case. You can add both targets to the same rule.

upvoted 1 times

bjexamprep 1 year, 5 months ago

Selected Answer: ACE

Step function is a process/workflow orchestrator. Usually process/workflow orchestrator doesn't do actual task, cause the objective of a orchestrator is to maintain the stage of a process/workflow. Instead, the orchestrator call a service to complete the task and update the stage.

So the task of removing access should be done by a Lambda function. Since lambda function is not an option, the only applicable option is C, while ECS introduces too much administration overhead, and is a very bad choice for this task.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: ADE

A, D and E

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: ADE

ADE. have to assume the step function calls lambda or some such to actually perform action

upvoted 1 times

Maria2023 2 years, 1 month ago

Selected Answer: ADE

I've chosen the EventBridge option (A) because I really was not able to find a way to set Cloudtrail to trigger SNS on it's own. The rest 2 are common sense

upvoted 2 times

AWSLord32 1 year, 5 months ago

Here you go https://docs.aws.amazon.com/awscloudtrail/latest/userguide/configure-sns-notifications-for-cloudtrail.html

upvoted 1 times

OCHT 2 years, 2 months ago

Selected Answer: ABE

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Define a pattern with the detail-type value set to AWS API Call via CloudTrail and an eventName of CreateUser.

B. Configure CloudTrail to send a notification for the CreateUser event to an Amazon Simple Notification Service (Amazon SNS) topic.

E. Use Amazon Simple Notification Service (Amazon SNS) to notify the security team.

upvoted 2 times

OCHT 2 years, 2 months ago

By creating an Amazon EventBridge rule, the company can detect the CreateUser event in CloudTrail and use it to trigger actions such as sending notifications or invoking AWS Lambda functions.

Configuring CloudTrail to send a notification for the CreateUser event to an Amazon SNS topic allows the security team to receive a notification whenever a new IAM user is created.

Using Amazon SNS, the security team can receive the notification and approve or deny the new IAM user creation. If the security team denies the creation, access can be automatically removed using AWS Lambda or AWS Step Functions.

Therefore, these three steps will allow the company to meet its requirements for user creation approval and access removal.

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: ADE

ADE is right

upvoted 1 times

[Removed] 2 years, 4 months ago

Selected Answer: ADE

ADE Step Functions works.

upvoted 1 times

Musk 2 years, 4 months ago

Selected Answer: ACE

I like ACE better. I am not sure Step Functions would work.

upvoted 1 times

moota 2 years, 4 months ago

According to ChatGPT, AWS Step Functions can interact with AWS APIs in a few different ways. One example is below.

Directly invoking AWS APIs using the "Task" state in Step Functions. This state type allows you to run an AWS Lambda function, which can interact with AWS APIs as part of its logic.

upvoted 1 times

zhangyu20000 2 years, 5 months ago

ADE are correct

upvoted 1 times

Question #118 *Topic 1*

A company wants to migrate to AWS. The company wants to use a multi-account structure with centrally managed access to all accounts and applications. The company also wants to keep the traffic on a private network. Multi-factor authentication (MFA) is required at login, and specific roles are assigned to user groups.

The company must create separate accounts for development. staging, production, and shared network. The production account and the shared network account must have connectivity to all accounts. The development account and the staging account must have access only to each other.

Which combination of steps should a solutions architect take 10 meet these requirements? (Choose three.)

A. Deploy a landing zone environment by using AWS Control Tower. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations.

B. Enable AWS Security Hub in all accounts to manage cross-account access. Collect findings through AWS CloudTrail to force MFA login.

C. Create transit gateways and transit gateway VPC attachments in each account. Configure appropriate route tables.

D. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.

E. Enable AWS Control Tower in all accounts to manage routing between accounts. Collect findings through AWS CloudTrail to force MFA login.

F. Create IAM users and groups. Configure MFA for all users. Set up Amazon Cognoto user pools and Identity pools to manage access to accounts and between accounts.

Suggested Answer: *BDF*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: ACD

The correct answer would be options A, C and D, because they address the requirements outlined in the question.

A. Deploying a landing zone environment using AWS Control Tower and enrolling accounts in an organization in AWS Organizations allows for a centralized management of access to all accounts and applications.

C. Creating transit gateways and transit gateway VPC attachments in each account and configuring appropriate route tables allows for private network traffic, and ensures that the production account and shared network account have connectivity to all accounts, while the development and staging accounts have access only to each other.

D. Setting up and enabling AWS IAM Identity Center (AWS Single Sign-On) and creating appropriate permission sets with required MFA for existing accounts allows for multi-factor authentication at login and specific roles to be assigned to user groups.

upvoted 17 times

masetromain 2 years, 5 months ago

The other options are not correct because:

B. Enabling AWS Security Hub in all accounts to manage cross-account access and collecting findings through AWS CloudTrail to force MFA login is not enough to meet the requirement of creating separate accounts for development, staging, production, and shared network. It can be used in addition to the other steps, but not as a standalone solution.

E. Enabling AWS Control Tower in all accounts to manage routing between accounts and collecting findings through AWS CloudTrail to force MFA login is not enough to meet the requirement of creating separate accounts for development, staging, production, and shared network. It can be used in addition to the other steps, but not as a standalone solution.

upvoted 4 times

masetromain 2 years, 5 months ago

F. Creating IAM users and groups and configuring MFA for all users and setting up Amazon Cognito user pools and Identity pools to manage access to accounts and between accounts does not address the requirement of creating separate accounts for development, staging, production, and shared network. Additionally, it does not address the requirement of keeping the traffic on a private network.

upvoted 3 times

amministrazione Most Recent 10 months ago

A. Deploy a landing zone environment by using AWS Control Tower. Enroll accounts and invite existing accounts into the resulting organization in AWS Organizations.

C. Create transit gateways and transit gateway VPC attachments in each account. Configure appropriate route tables.

D. Set up and enable AWS IAM Identity Center (AWS Single Sign-On). Create appropriate permission sets with required MFA for existing accounts.

upvoted 1 times

ajeeshb 1 year, 3 months ago

Selected Answer: ACD

A, C and D are right answers. Option C is though not clear. Transit gateway needs to be created in shared network account and tgw vpc attachment in all accounts. But option C says "create tgw and tgw vpc attachment in all accounts", which is a bit confusing

upvoted 2 times

8693a49 11 months ago

Yes, you probably only need one TGW in the shared account

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: ACD

A, C and D

upvoted 1 times

shaaam80 1 year, 6 months ago

Selected Answer: ACD

Answer - ACD

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: ACD

ACD easy

upvoted 1 times

Maria2023 2 years ago

Selected Answer: ACD

ACD seems like the only technically achievable solution. B and E appear to be completely wrong and for F - I am not sure whether Cognito will do the job but for sure it would be extremely hard to implement that way.

upvoted 2 times

OCHT 2 years, 2 months ago

Selected Answer: ACD

Option E is not the most appropriate choice because it suggests enabling AWS Control Tower in all accounts to manage routing between accounts. However, AWS Control Tower is not primarily designed for managing routing between accounts; it is intended to set up and govern a secure, multi-account AWS environment. The transit gateways and VPC attachments in Option C are better suited for managing routing and connectivity between accounts.

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: ACD

ACD are the best choice

upvoted 1 times

spd 2 years, 4 months ago

Selected Answer: ACD

By Elimination Rule

upvoted 3 times

zhangyu20000 2 years, 5 months ago

ACD are correct.

upvoted 3 times

Question #119 *Topic 1*

A company runs its application in the eu-west-1 Region and has one account for each of its environments: development, testing, and production. All the environments are running 24 hours a day, 7 days a week by using stateful Amazon EC2 instances and Amazon RDS for MySQL databases. The databases are between 500 GB and 800 GB in size.

The development team and testing team work on business days during business hours, but the production environment operates 24 hours a day, 7 days a week. The company wants to reduce costs. All resources are tagged with an environment tag with either development, testing, or production as the key.

What should a solutions architect do to reduce costs with the LEAST operational effort?

A. Create an Amazon EventBridge rule that runs once every day. Configure the rule to invoke one AWS Lambda function that starts or slops instances based on me tag, day, and time.

B. Create an Amazon EventBridge rule that runs every business day in the evening. Configure the rule to invoke an AWS Lambda function that stops instances based on the tag. Create a second EventBridge rule that runs every business day in the morning. Configure the second rule lo invoke another Lambda function that starts instances based on the tag.

C. Create an Amazon EventBridge rule that runs every business day in the evening, Configure the rule to invoke an AWS Lambda function that terminates, instances based on the lag. Create a second EventBridge rule that runs every business day in the morning. Configure the second rule lo invoke another Lambda function that restores the instances from their last backup based on the tag.

D. Create an Amazon EventBridge rule that runs every hour. Configure the rule to invoke one AWS Lambda function that terminates or restores instances from their last backup based on the tag. day, and time.

Suggested Answer: *A*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: B

The correct answer is B. Creating an Amazon EventBridge rule that runs every business day in the evening to stop instances and another rule that runs every business day in the morning to start instances based on the tag will reduce costs with the least operational effort.

This approach allows for instances to be stopped during non-business hours when they are not in use, reducing the costs associated with running them. It also allows for instances to be started again in the morning when the development and testing teams need to use them.

Option A would require the instances to be stopped and started once a day, which could result in instances being stopped while they are in use or not being stopped when they are not in use.

Option C would terminate instances during non-business hours and restore them again in the morning, which could lead to data loss or longer start up times.

Option D would terminate or restore instances every hour, which could lead to unnecessary costs as well as data loss or longer start up times.

upvoted 11 times

Musk Highly Voted 2 years, 4 months ago

Selected Answer: B

this is easy. I wish I'll have several of this in the exam.

upvoted 8 times

nimbus_00 Most Recent 8 months, 3 weeks ago

Selected Answer: B

Stopping instances rather than terminating them ensures that the environment's state can be quickly restored the next day without needing to manage backups or restorations, making it operationally efficient.

upvoted 1 times

amministrazione 10 months ago

B. Create an Amazon EventBridge rule that runs every business day in the evening. Configure the rule to invoke an AWS Lambda function that stops instances based on the tag. Create a second EventBridge rule that runs every business day in the morning. Configure the second rule lo invoke another Lambda function that starts instances based on the tag.

upvoted 1 times

AWSLord32 1 year, 5 months ago

Selected Answer: B

Voted B, but C seems to be more cost effective. Any idea to why it wouldn't work?

upvoted 1 times

pangchn 1 year, 4 months ago

C will terminate the instance which may potentially the work on the disk

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: B

Option B

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

B for sure

upvoted 1 times

Maria2023 2 years ago

Selected Answer: B

A cannot complete the requirement since it runs once a day and we need to stop the non-prod instances in the eveninig and start them in the morning. A would potentially work if we set up the rule to run every hour and then determine the appropriate action based on the time of the day. C and D are nonsense to me

upvoted 1 times

leehjworking 2 years, 1 month ago

Can anyone explain why B has less operational effort than A

upvoted 1 times

chikorita 2 years ago

cuz we have to schedule Eventbridge to run twice a day [STOP trigger and START trigger]....Option A mentions about "ONCE" which could only be either stop or start so option B is most appropiate

upvoted 1 times

dev112233xx 2 years, 2 months ago

Selected Answer: B

B is correct

The keyword here is whether you terminate or stop the instance. ofc you don't want to terminate. stop is enough and company don't pay when the instance is in stop state.

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: B

B is the easy choice

upvoted 2 times

zhangyu20000 2 years, 5 months ago

B is correct. Stop the instance that preserver all data.

C: is incorrect because it terminate instance that will loss data

upvoted 5 times

rbm2023 2 years, 1 month ago

with the addition to the fact that to recreate those DBs from scratch would take a long time.

upvoted 2 times

Question #120 *Topic 1*

A company is building a software-as-a-service (SaaS) solution on AWS. The company has deployed an Amazon API Gateway REST API with AWS Lambda integration in multiple AWS Regions and in the same production account.

The company offers tiered pricing that gives customers the ability to pay for the capacity to make a certain number of API calls per second. The premium tier offers up to 3,000 calls per second, and customers are identified by a unique API key. Several premium tier customers in various Regions report that they receive error responses of 429 Too Many Requests from multiple API methods during peak usage hours. Logs indicate that the Lambda function is never invoked.

What could be the cause of the error messages for these customers?

A. The Lambda function reached its concurrency limit.

B. The Lambda function its Region limit for concurrency.

C. The company reached its API Gateway account limit for calls per second.

D. The company reached its API Gateway default per-method limit for calls per second.

Suggested Answer: *C*

Community vote distribution

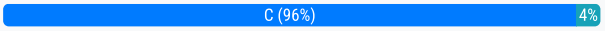

sambb Highly Voted 2 years, 3 months ago

Selected Answer: C

API Gateway has a limit of 10k requests per second, per account, per region
https://docs.aws.amazon.com/apigateway/latest/developerguide/limits.html
upvoted 13 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct answer is C. The company reached its API Gateway account limit for calls per second. This is because Amazon API Gateway has a default account-level limit of 10,000 requests per second (RPS) and a default per-method limit of 5,000 RPS. If the company's premium tier customers are making more than 10,000 requests per second in total across all API methods and regions, they would be receiving the error message of 429 Too Many Requests. This indicates that the API Gateway account is reaching its capacity limit, and the Lambda function is not being invoked because API Gateway is blocking the requests before they reach the Lambda function.

The other choices are not correct because the Lambda function's concurrency limit and region limit for concurrency would not affect the API Gateway's request rate limit, and the API Gateway's default per-method limit is 5,000 RPS which is less than the premium tier's 3,000 calls per second.
upvoted 6 times

masetromain 2 years, 5 months ago

Option A is incorrect because the error message is not related to the Lambda function reaching its concurrency limit.

Option B is incorrect because the error message is not related to the Lambda function reaching its region limit for concurrency.

Option D is incorrect because the error message is not related to the company reaching its API Gateway default per-method limit for calls per second, but it's related to the account level limit.
upvoted 4 times

Paul123456789 Most Recent 2 months, 4 weeks ago

Selected Answer: D

"they receive error responses of 429 Too Many Requests from multiple API methods"
Error are not coming from all API calls but from multiple API methods
D looks correct
upvoted 1 times

BennyMao 3 months, 2 weeks ago

Selected Answer: D

API Gateway enforces a default per-method limit of 1,000 RPS (requests per second).
Since premium customers require up to 3,000 RPS, they would exceed this limit, leading to 429 errors.

upvoted 1 times

E90 5 months ago

Selected Answer: D

Going with D because:
1) 429 "Too Many requests" error is shown
"Check the rate or burst limit for per-client or per-method throttling limits that you set for the API stage for your usage plan. When the rate or burst limit is exceeded, the CloudWatch event logs an exceeded throttle limit."
https://repost.aws/knowledge-center/api-gateway-429-limit

2) The premium tier offer only allows up to 3000 calls per second, while API Gateway has a limit of 10k requests per second, per account, per region - this is far below the allocated limit

3) Question also points to "multiple premium tier users in various regions" i.e. this issue is happening for different AWS accounts in diferrent regions - which suggests that it is not related to the account limit

upvoted 1 times

sintesi_suffisso0 5 months ago

Actually the account is only one

upvoted 2 times

Heman31in 6 months, 2 weeks ago

Selected Answer: D

The most likely cause of the 429 Too Many Requests error messages, despite the Lambda function not being invoked, is D. The company reached its API Gateway default per-method limit for calls per second.

Here's a breakdown of why:

API Gateway Limits: API Gateway imposes rate limits on API methods to prevent abuse and ensure fair resource allocation. These limits can be configured at the account level or the individual method level.
Default Limits: If not explicitly configured, API Gateway applies default limits to methods. These default limits may be insufficient for high-traffic scenarios, especially during peak usage hours.
Lambda Function Invocations: The Lambda function is not invoked because the request is being throttled at the API Gateway level before it reaches the Lambda function.

upvoted 1 times

amministrazione 10 months ago

C. The company reached its API Gateway account limit for calls per second.

upvoted 1 times

8693a49 11 months ago

Selected Answer: B

I'm going to argue the problem is that the source of the errors is Lambda reaching it's regional concurency limit.

By default, Lambda has a regional limit of 1000 concurrent invocations. The premium tier allows 3000 requests/s. Depending on the number of premium customers and the average duration of a call it may be that the concurency limit is reached or not. We don't know for sure, but it is certainly plausible. If Lambda hits the concurrency limit, it also returns 429. So how do we know where the error is coming from?

The key is in what exactly fails: "multiple API methods during peak usage hours". Notice it is not ALL API calls. If the gateway was rate limiting then we would see blocks of random requests being denied until the rate bucket empties to allow new ones. But if the Lambda concurrency is hit then it means no new execution environments of Lambda are created, but the ones that exist keep processing. So the behaviour is that some functions will continue to operate, while others will start thorttling with 429. This behaviour better matches with "multiple API methods" failing.

upvoted 1 times

helloworldabc 10 months ago

just C

upvoted 1 times

fangd0n 1 year, 1 month ago

C correct. This is Gateway API response

upvoted 1 times

JohnLuo 1 year, 2 months ago

Selected Answer: C

C is correct.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

429 is API Gateway API throttle default limit.

upvoted 3 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C of course

upvoted 1 times

dev112233xx 2 years, 2 months ago

Selected Answer: C

C

429 error indicates that API calls per second was exceeded ... it's not a Lambda issue

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: C

Company reached its limit

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: C

C is the answer

upvoted 1 times

zhangyu20000 2 years, 5 months ago

C is correct answer

upvoted 1 times

Question #121 *Topic 1*

A financial company is planning to migrate its web application from on premises to AWS. The company uses a third-party security tool to monitor the inbound traffic to the application. The company has used the security tool for the last 15 years, and the tool has no cloud solutions available from its vendor. The company's security team is concerned about how to integrate the security tool with AWS technology.

The company plans to deploy the application migration to AWS on Amazon EC2 instances. The EC2 instances will run in an Auto Scaling group in a dedicated VPC. The company needs to use the security tool to inspect all packets that come in and out of the VPC. This inspection must occur in real time and must not affect the application's performance. A solutions architect must design a target architecture on AWS that is highly available within an AWS Region.

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

A. Deploy the security tool on EC2 instances m a new Auto Scaling group in the existing VPC

B. Deploy the web application behind a Network Load Balancer

C. Deploy an Application Load Balancer in front of the security tool instances

D. Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool

E. Provision a transit gateway to facilitate communication between VPCs.

Suggested Answer: *AD*

Community vote distribution

AD (49%) DE (49%)

rbm2023 Highly Voted 2 years, 1 month ago

Selected Answer: DE

Based on the scenario in question, the requirement is that the security tool will run in an auto scaling group in a dedicated VPC this cannot be changed. This will break Option A. If we look at the usage for the Gateway Load Balancer which is the key for the solution where application cannot have performance hits if you are inspecting the traffic, so you need to TAP the traffic to move into another third-party tool. In the references you will find below the transit gateway will facilitate the VPC-to-VPC communication and as you can see, the security appliances VPC is a segregated from the application VPC, so again, option A is NOT valid.
https://catalog.workshops.aws/networking/en-US/gwlb
https://www.fortinet.com/blog/business-and-technology/highly-scalable-fortigate-next-generation-firewall-security-on-aws-gateway-load-balancer-service
upvoted 24 times

OCHT Highly Voted 2 years, 2 months ago

Selected Answer: AD

Option B, deploying the web application behind a Network Load Balancer, is not relevant to integrating the third-party security tool with AWS technology.

Option C, deploying an Application Load Balancer in front of the security tool instances, is not necessary because a Gateway Load Balancer is already being used to redirect traffic to the security tool.

Option E, provisioning a transit gateway to facilitate communication between VPCs, is not relevant to integrating the third-party security tool with AWS technology or inspecting packets in and out of the VPC.

In summary, options A and D are the best choices because address the specific requirements stated in the scenario while options B, C and E do not.
upvoted 22 times

43c89f4 1 year, 2 months ago

DE is correct, the question clearly mention which combination
- GWLB and provision transit gateway is solution
upvoted 3 times

deegadaze1 2 years, 1 month ago

Correct for GLB---> https://www.youtube.com/watch?v=-j2smz_VCH4

upvoted 2 times

Kaps443 Most Recent 2 weeks, 6 days ago

Selected Answer: AD

A. Deploy the security tool on EC2 instances
Since the tool is legacy and has no cloud-native support, it must run on EC2 instances.

An Auto Scaling group helps maintain availability and resilience.

Placing the tool in the same VPC allows for easy traffic routing via a Gateway Load Balancer.

D. Provision a Gateway Load Balancer (GWLB) per AZ
Gateway Load Balancer is designed specifically for this use case: deploying virtual appliances (like firewalls, IDS/IPS, and packet inspection tools).

It redirects VPC traffic to your EC2-based security tool instances for inline, transparent inspection, without affecting the app performance.

upvoted 1 times

jimee11 1 month, 3 weeks ago

Selected Answer: AD

Requirements clearly call out the 'web application' vs. 'security tool'. The web application is not going to be deployed behind an NLB. This rules out B.
Requirements say you need scalability. Deploy the security tool behind an ALB with auto-scaling (A). Gateway LB is best for deep packet inspections (D).

upvoted 1 times

kylix75 5 months ago

Selected Answer: BD

Correct answers: B and D

Rationale:
B (Network Load Balancer):
- Operates at layer 4
- Minimal latency impact
- Supports transparent network inspection

D (Gateway Load Balancer):
- Purpose-built for third-party security appliances
- Enables inline traffic inspection
- High availability with per-AZ deployment

Other options' issues:
A: Doesn't address traffic routing
C: ALB operates at layer 7, adding unnecessary overhead
E: Transit gateway unnecessary for single VPC setup

upvoted 1 times

ahhatem 6 months, 2 weeks ago

Selected Answer: DE

The question explicitly states that it would be deployed in a dedicated VPC. This disqualifies A.
On another hand, dedicated security appliances are usually deployed in a centralized networking setup with central ingress/egress. Check this whitepaper:

https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/using-gwlb-with-tg-for-cns.html

upvoted 1 times

Edd_18 7 months, 1 week ago

Selected Answer: DE

https://www.fortinet.com/blog/business-and-technology/highly-scalable-fortigate-next-generation-firewall-security-on-aws-gateway-load-balancer-service

upvoted 2 times

FZA24 8 months, 2 weeks ago

Selected Answer: AD

In AD, it mention that will be deployed in the existing VPC. however, in DE, it does not mention that the security tool is deployed in another VPC. It only mention transit gateway between VPCs.

upvoted 2 times

AWSum1 8 months, 4 weeks ago

Selected Answer: AD

it says it needs to inspect traffic coming in and out of THE VPC not multiple VPC's. This statement disqualifies E

upvoted 3 times

amministrazione 10 months ago

A. Deploy the security tool on EC2 instances m a new Auto Scaling group in the existing VPC

D. Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool

upvoted 1 times

ry1999 10 months ago

Selected Answer: AD

D and E make the most sense if your architecture involves multiple VPCs where traffic needs to be centrally managed and inspected. This combination addresses both the direct need for packet inspection and the broader network management requirements.

A and E could be considered if the application and security tool deployment are straightforward and confined to a single or connected VPCs. However, managing traffic flow effectively to the security tools might require additional configuration that can complicate the setup.

Since there is only once VPC, AD

upvoted 3 times

Jason666888 10 months, 4 weeks ago

Selected Answer: AD

It has to be AD.

I've taken the Udemy Course from stephane maarek and his course described this kind of scenario

upvoted 5 times

seochan 11 months ago

Selected Answer: AD

DE cannot be the answer. The combination doesn't describe how to deploy the security tools on the cloud.

upvoted 2 times

michele_scar 1 year ago

Selected Answer: DE

DE is the answer. Transit -> GWLB -> Inspection tool

upvoted 1 times

helloworldabc 10 months ago

just AD

upvoted 1 times

ce825d4 1 year ago

Selected Answer: AD

AD is correct as the requirement is to use the Security tool to inspect traffic coming in and out of the VPC. So, you need to deploy the security tool on EC2 instances and provision a Gateway loadbalancer to load balance the traffic. With a GLB, you can deploy, manage, and scale virtual appliances, such as intrusion detection and prevention, firewalls, and deep packet inspection systems. It creates a single entry and exit point for all appliance traffic and scales your virtual appliances with demand. You can also exchange traffic across virtual private cloud (VPC) boundaries.

upvoted 2 times

seetpt 1 year, 1 month ago

Selected Answer: AD

AD for me

upvoted 2 times

red_panda 1 year, 2 months ago

Selected Answer: DE

DE without doubts guys.

GLB is just for this reason. Deploy the security tool into another ASG will only increase the cost and it's crazy, the performance isn't the same as the GLB (which operates at Lv. 3 of networking).

upvoted 2 times

Question #122 *Topic 1*

A company has purchased appliances from different vendors. The appliances all have IoT sensors. The sensors send status information in the vendors' proprietary formats to a legacy application that parses the information into JSON. The parsing is simple, but each vendor has a unique format. Once daily, the application parses all the JSON records and stores the records in a relational database for analysis.

The company needs to design a new data analysis solution that can deliver faster and optimize costs.

Which solution will meet these requirements?

A. Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon. S3 Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis.

B. Migrate the application server to AWS Fargate, which will receive the information from IoT sensors and parse the information into a relational format. Save the parsed information to Amazon Redshlft for analysis.

C. Create an AWS Transfer for SFTP server. Update the IoT sensor code to send the information as a .csv file through SFTP to the server. Use AWS Glue to catalog the files. Use Amazon Athena for analysis.

D. Use AWS Snowball Edge to collect data from the IoT sensors directly to perform local analysis. Periodically collect the data into Amazon Redshift to perform global analysis.

Suggested Answer: *C*

Community vote distribution

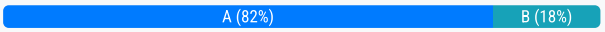

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: A

IOT Core communication supports protocols MQTT, HTTPS, MQTT over WSS, and LoRaWAN (but not FTP/SFTP) so C should be wrong.

Rules Engine: AWS IoT Core provides a rules engine that allows users to define and execute business logic on the data generated by their IoT devices. This enables users to automate actions such as sending notifications, triggering alarms, or updating device settings based on real-time data.

Integration with other AWS Services: AWS IoT Core integrates with other AWS services such as AWS Lambda, AWS Kinesis, and AWS S3, allowing users to easily process and store their IoT data, as well as build complex IoT applications using a range of AWS services.

upvoted 12 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: A

A. Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon S3. Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis.

This solution meets the requirement of faster analysis and cost optimization by using AWS IoT Core to collect data from the IoT sensors in real-time and then using AWS Glue and Amazon Athena for efficient data analysis.

Option B and D do not optimize the cost of data analysis as they involve use of expensive services like AWS Fargate and Snowball Edge respectively. Option C does not make use of real-time data collection and may not be optimal for faster analysis.

upvoted 5 times

amministrazione Most Recent 10 months ago

A. Connect the IoT sensors to AWS IoT Core. Set a rule to invoke an AWS Lambda function to parse the information and save a .csv file to Amazon. S3 Use AWS Glue to catalog the files. Use Amazon Athena and Amazon QuickSight for analysis.

upvoted 1 times

gofavad926 1 year, 3 months ago

Selected Answer: A

A, IoT Core

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: A

Option A is best(fastest) and most cost effective.

upvoted 1 times

GaryQian 1 year, 7 months ago

Selected Answer: A

Everytime the exam shows IOT sensor, think of IOT Core and aws glue

upvoted 1 times

KCjoe 1 year, 8 months ago

Selected Answer: B

How can A satisfy this requirement? "relational database for analysis"

The only option is B with relational database for analysis.

upvoted 4 times

helloworldabc 10 months ago

just A

upvoted 1 times

heatblur 1 year, 7 months ago

"The company needs to design a new data analysis solution that can deliver faster and optimize costs."

upvoted 2 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: A

Option A: AWS IoT Core + Lambda

Speed: Near real-time data collection and analysis.

Flexibility: Ability to adapt to different data formats from multiple vendors.

Option C: AWS Transfer for SFTP

Speed: There may be network delays and waiting for all data to be sent.

Development needs: The sensor code needs to be updated, which increases the development workload.

All things considered, option A is better than option C in terms of speed and flexibility, and is especially suitable for real-time or near-real-time requirements.

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A for sure

upvoted 1 times

Maria2023 2 years, 1 month ago

Selected Answer: A

I go for A on the elimination principle although neither of the answers does not seem to fully cover the requirements. I am not sure what is the "vendors' proprietary formats" and not sure why they assume it's csv. Also there is a requirement to load the data in relational database which excludes B. For A we need to assume that S3 covers this requirement.

upvoted 2 times

dev112233xx 2 years, 2 months ago

Selected Answer: A

A is correct, even though it's not clear from the question if the sensors protocol is MQTT or HTTPS.

but i can't find other suitable answer so i guess A is the correct one.

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: A

Connect the IoT sensors to AWS IoT Core.

upvoted 2 times

spd 2 years, 4 months ago

Selected Answer: A

A by Elimination rule

upvoted 4 times

Musk 2 years, 4 months ago

Selected Answer: B

I m not convinced about A. It kind of requires changes in the sensors to be compatible with AWS IoT Core.

upvoted 4 times

Sarutobi 2 years, 2 months ago

I agree with you here. We don't know if IoT Core supports it, so moving the application to AWS Fargate will guarantee compatibility.

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: A

i'll go for A

upvoted 4 times

zhangyu20000 2 years, 5 months ago

A is correct.

B: it is appliance, impossible to install on Fargate

C: device not use FTP protocol

D: snowball is not real time

upvoted 4 times

Musk 2 years, 4 months ago

In B, we don't try to port appliances to Fargate, but only the app that parses the informtion from the appliances into JSON.

I am doubting about A. Unless you would reprogrm the sensors they would not know how to connect to AWS IoT Core.

upvoted 1 times

Question #123 *Topic 1*

A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connect connection in a central network account.

The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.

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

A. Create a Direct Connect gateway in the central account. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.

B. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.

C. Provision an internet gateway. Attach the internet gateway to subnets. Allow internet traffic through the gateway.

D. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.

E. Provision VPC peering as necessary.

F. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

Suggested Answer: *BDF*

Community vote distribution

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: BDF

B and D and F are correct.

B: Creating a Direct Connect gateway and a transit gateway in the central network account will allow the company to connect its on-premises data center to the resources in AWS.

D: Sharing the transit gateway with other accounts will allow the company to communicate with all the VPCs in multiple accounts.

F: Provisioning only private subnets and opening necessary routes on the transit gateway and customer gateway will allow the company to route its cloud resources to the internet through its on-premises data center.

A is incorrect because it would be redundant to use both a Direct Connect gateway and a transit gateway.

C is incorrect because it is not necessary to provision an internet gateway, since the company wants to route traffic through their on-premises data center.

E is incorrect because VPC peering may not be necessary if the company is using a transit gateway to connect all the VPCs.

upvoted 13 times

amministrazione Most Recent 10 months ago

B. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.

D. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.

F. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: BDF

BDF is most scalable solution.

upvoted 1 times

shaaam80 1 year, 6 months ago

Selected Answer: BDF

Answer BDF

DGW and TGW

Share TGW and configure VPC attachments to TGW

Open necessary routes for traffic routing via NAT gw on the on-prem dc

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: BDF

Very logical

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: BDF

BDF for sure

upvoted 1 times

Maria2023 2 years ago

Selected Answer: BDF

Standard scenario. You connect the Direct Connect Gateway to the Transit Gateway, attach the VPCs, and route the traffic through the On-premise devices

upvoted 3 times

SkyZeroZx 2 years, 1 month ago

Selected Answer: BDF

BDF is the right ans

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: BDF

BDF is the right combo

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: BDF

VPC Peering does not work as there are hundreds of VPCs, transit gateway is easy to configure and practical.
https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html

upvoted 4 times

zozza2023 2 years, 5 months ago

Selected Answer: BDF

B D and F

upvoted 4 times

zozza2023 2 years, 5 months ago

I agree with BD&F

upvoted 3 times

zhangyu20000 2 years, 5 months ago

BDF are correct

upvoted 2 times

Question #124 *Topic 1*

A company has hundreds of AWS accounts. The company recently implemented a centralized internal process for purchasing new Reserved Instances and modifying existing Reserved Instances. This process requires all business units that want to purchase or modify Reserved Instances to submit requests to a dedicated team for procurement. Previously, business units directly purchased or modified Reserved Instances in their own respective AWS accounts autonomously.

A solutions architect needs to enforce the new process in the most secure way possible.

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

A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.

B. Use AWS Config to report on the attachment of an IAM policy that denies access to the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.

C. In each AWS account, create an IAM policy that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action.

D. Create an SCP that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action. Attach the SCP to each OU of the organization.

E. Ensure that all AWS accounts are part of an organization in AWS Organizations that uses the consolidated billing feature.

Suggested Answer: *AD*

Community vote distribution

AD (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: AD

A and D are the correct answer.
A: By ensuring all AWS accounts are part of an organization in AWS Organizations, it allows for centralized management and control of the accounts. This can help enforce the new purchasing process by giving a dedicated team the ability to manage and enforce policies across all accounts.
D: By creating an SCP (Service Control Policy) that denies access to the ec2:PurchaseReservedInstancesOffering and ec2:ModifyReservedInstances actions, it enforces the new centralized purchasing process. Attaching the SCP to each OU (organizational unit) within the organization ensures that all business units are adhering to the new process.

B and C are not the correct answer, because AWS Config and IAM policies are used for monitoring and managing access to resources in an account, respectively. They don't enforce the new process for purchasing reserved instances.
E is not the correct answer as this is not related to the new process for purchasing reserved instances.

upvoted 9 times

amministrazione Most Recent 10 months ago

A. Ensure that all AWS accounts are part of an organization in AWS Organizations with all features enabled.
D. Create an SCP that denies the ec2:PurchaseReservedInstancesOffering action and the ec2:ModifyReservedInstances action. Attach the SCP to each OU of the organization.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: AD

A and D

upvoted 1 times

atirado 1 year, 6 months ago

A+D achieve the goal of denying access to purchase and to modify Reserved Instances to all OUs. The dedicated team can still perform these actions if they are part of the management account.

C, E don't actually do anything, as in, actually control anything at all. B will trigger on the wrong thing to be alarmed about, if triggering an alarm was the goal.

upvoted 1 times

dkcloudguru 1 year, 9 months ago

A and D : is the best way

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: AD

AD. A so can use SCP

upvoted 1 times

Maria2023 2 years, 1 month ago

Selected Answer: AD

I was not confident about enabling all features because I was messing "features" and "services". Yes - you need to enable all features, otherwise you cannot control the accounts in your organization. The rest is common sense

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: AD

AD easy

upvoted 3 times

zozza2023 2 years, 5 months ago

Selected Answer: AD

A and D

upvoted 4 times

zhangyu20000 2 years, 5 months ago

AD are correct

upvoted 2 times

Question #125 *Topic 1*

A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.

A recent RDS database failover test caused a 40-second outage to the application. A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.

Which combination of steps should the solutions architect take to meet these requirements? (Choose three.)

A. Use Amazon ElastiCache for Memcached in front of the database

B. Use Amazon ElastiCache for Redis in front of the database

C. Use RDS Proxy in front of the database.

D. Migrate the database to Amazon Aurora MySQL.

E. Create an Amazon Aurora Replica.

F. Create an RDS for MySQL read replica

Suggested Answer: *BCF*

Community vote distribution

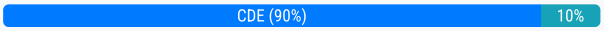

RaghavendraPrakash Highly Voted 2 years, 1 month ago

CDE. RDS Failover typically takes 60-120 seconds, while Aurora failover completes within 30 seconds. ElastiCache is for reducing latency, not for failover.

upvoted 13 times

dev112233xx Highly Voted 2 years, 2 months ago

Selected Answer: CDE

RDS Proxy with Aurora are the best combination for less than "20 sec" failover time...

According to this article RDS Proxy can reduce the failover time of Aurora by 79% while it can reduce RDS failover time by only 32%:

https://aws.amazon.com/blogs/database/improving-application-availability-with-amazon-rds-proxy/

upvoted 10 times

zak543 Most Recent 5 months ago

Selected Answer: CDE

for multiple chose answer, if i select 2of3 right, in this case the whole quest will be wrong or what?

upvoted 1 times

amministrazione 10 months ago

C. Use RDS Proxy in front of the database.

D. Migrate the database to Amazon Aurora MySQL.

E. Create an Amazon Aurora Replica.

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: CDE

A and B don't contribute to reducing response time in failover scenarios.

D is required for faster failover.

E is required to support D.

F doesn't reduce failover time.

C, finally, is the remaining option. It doesn't hurt, and can contribute to faster failover, though it is not the most important factor here - the switch to Aurora with an Aurora read replica is.

upvoted 2 times

bjexamprep 1 year, 3 months ago

Anyone can share why choosing E? I know we have to choose 3. isn't it weird? Aurora already has replicas natively. Why creating another one?

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: CDE

Option C, D and E

upvoted 1 times

atirado 1 year, 6 months ago

Selected Answer: CDE

C+D+E provides the 'fastest' failover with the options available:

- Aurora MySQL is Multi-AZ by design: During failover it will promote a Replica to primary or create a Primary instance
- Creating a Replica provides the option to have something to failover to
- Using an RDS Proxy further reduces failover time and provides 'transparent' failovers as well (It manages DNS changes)

The argument against Caching (options A or B) is that it doesn't accelerate failing over to a different instance. Cache misses and write operations will produce exceptions because there is no instance to query. Moreover, there is no information in the question to choose between either caching option, i.e. Both options can be created starting from an Aurora DB Cluster settings - https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/creating-elasticache-cluster-with-RDS-settings.html

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: CDE

CDE, agree with other comments

upvoted 2 times

Sarutobi 2 years, 2 months ago

Selected Answer: CDE

The trick seems to be that the RDS proxy handles DNS updates quickly. While if you don't use it, you are at the mercy of the host to update its DNS cache.

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: CDE

CDE is the best choice

upvoted 1 times

DWsk 2 years, 3 months ago

Selected Answer: CDE

CDE. I would have said F, but the question asks for a combination of steps, so its looking for the Aurora replica and not the MySQL RDS replica

upvoted 3 times

Jay_2pt0_1 2 years, 2 months ago

I agree with your logic.

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: CDE

C for sure as connection pooling helps quick re connect. There is no preference for A or B cache solution based on the question. So, A,B are eliminated. so three correct options should be in others. If you choose Aurora only, three answers will be met :-) C,D,E

upvoted 3 times

zozza2023 2 years, 5 months ago

Selected Answer: CDE

C D and E

upvoted 2 times

nyxs_19 2 years, 4 months ago

A and B are incorrect options because Amazon ElastiCache is a caching service, not a failover solution. F is also incorrect because RDS read replicas are asynchronous, which means that there may be a delay in replication, leading to the potential loss of data. Additionally, creating a read replica does not improve the failover time.

upvoted 2 times

AjayD123 2 years, 5 months ago

Selected Answer: CDE

RDS read replica auto failover takes approx 35 seconds hence, BCF does not satisfy under 20 seconds failover requirement.
https://aws.amazon.com/rds/features/multi-az/#:~:text=Amazon%20RDS%20Multi%2DAZ%20with%20two%20readable%20standbys,-Automatically%20fail%20over&text=Automatically%20failover%20in%20typically%20under,and%20with%20no%20manual%20intervention.
upvoted 5 times

zozza2023 2 years, 5 months ago
thanks for the information about RDS read replica
upvoted 2 times

masetromain 2 years, 5 months ago
Selected Answer: CDE
The correct answer is D, E and C:

Migrate the database to Amazon Aurora MySQL.
- Create an Amazon Aurora Replica.
- Use RDS Proxy in front of the database.
- These options are correct because they address the requirement of reducing the failover time to less than 20 seconds.

Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called "Aurora Global Database" which allows you to create read-only replicas across multiple AWS regions which can further help to reduce the failover time.

Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure.

Using RDS proxy can also help to reduce the failover time as it can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.
upvoted 4 times

masetromain 2 years, 5 months ago
Option A and B, Use Amazon ElastiCache for Memcached and Redis in front of the database, are not correct as ElastiCache is a caching service, it doesn't provide a high availability solution for the underlying database.

Option F, Create an RDS for MySQL read replica, is not correct as a read replica can only be used to offload read traffic from the primary instance, it doesn't provide a high availability solution for the underlying database.
upvoted 1 times

masetromain 2 years, 5 months ago
Selected Answer: BCF
The correct answer is B, C and F.

Using Amazon ElastiCache for Redis in front of the database (Option B) will help to reduce the failover time by caching the frequently-used data, so that it can be quickly served from the cache rather than having to be retrieved from the database during a failover.

Using RDS Proxy in front of the database (Option C) will help to reduce the failover time by managing the connections to the RDS DB instance, so that it can quickly route traffic to the new primary instance during a failover.

Creating an RDS for MySQL read replica (Option F) will help to reduce the failover time by having a read-only copy of the database running in parallel with the primary instance, so that it can take over as the primary instance in the event of a failover.

Option A and D are not relevant in this case as the question is asking specifically about reducing failover time for an RDS for MySQL database.
upvoted 4 times

spd 2 years, 4 months ago
C, D and E Correct
upvoted 2 times

Question #126 *Topic 1*

An AWS partner company is building a service in AWS Organizations using its organization named org1. This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2. The company must establish least privilege security access using an API or command line tool to the customer account.

What is the MOST secure way to allow org1 to access resources in org2?

A. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks.

B. The customer should create an IAM user and assign the required permissions to the IAM user. The customer should then provide the credentials to the partner company to log in and perform the required tasks.

C. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) when requesting access to perform the required tasks.

D. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN), including the external ID in the IAM role's trust policy, when requesting access to perform the required tasks.

Suggested Answer: *D*

Community vote distribution

D (100%)

dev112233xx Highly Voted 2 years, 2 months ago

Selected Answer: D

D

Well.. "external ID" is the keyword that you should look for in such scenario.

upvoted 6 times

d401c0d Most Recent 5 months ago

Selected Answer: D

D. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN), including the external ID in the IAM role's trust policy, when requesting access to perform the required tasks.

A - that is just hilarious and should not be the case.

upvoted 1 times

amministrazione 10 months ago

D. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN), including the external ID in the IAM role's trust policy, when requesting access to perform the required tasks.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: D

Option D is most secure.

upvoted 1 times

atirado 1 year, 6 months ago

Selected Answer: D

Sharing credentials will always be a bad idea. In comparison to C and D, options A and B are insecure.

The reason D is the most secure option compared to C is because it addresses the confused deputy problem - https://docs.aws.amazon.com/IAM/latest/UserGuide/confused-deputy.html

upvoted 3 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

it's D, but private link would be a better choice

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: D

With the external ID.

upvoted 2 times

God_Is_Love 2 years, 3 months ago

Selected Answer: D

{
"Version": "2012-10-17",
"Statement": {
"Effect": "Allow",
"Principal": {
"AWS": "Example Corp's AWS Account ID"
},
"Action": "sts:AssumeRole",
"Condition": {
"StringEquals": {
"sts:ExternalId": "1122334455-The ID that only Third party and customer knows"
}
}
}
}

upvoted 3 times

Musk 2 years, 4 months ago

Selected Answer: D

Easy. The external ID is for sure the winner.

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: D

D seems the correct answer

upvoted 2 times

Untamables 2 years, 5 months ago

Selected Answer: D

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_third-party.html

upvoted 2 times

masetromain 2 years, 5 months ago

Selected Answer: D

The correct answer is D. This is the most secure way to allow org1 to access resources in org2 because it allows for least privilege security access. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) and include the external ID in the IAM role's trust policy when requesting access to perform the required tasks. This ensures that the partner company can only access the resources that it needs and only from the specific customer account.

Option A and B both involve providing the partner company with credentials, which can be easily compromised and could lead to a security breach. Option C also provides the partner company with an IAM role, but it doesn't have any restrictions on when and where the partner company can access the resources in customer account, it could be a security risk.

upvoted 3 times

zhangyu20000 2 years, 5 months ago

D is correct

upvoted 1 times

Question #127 *Topic 1*

A delivery company needs to migrate its third-party route planning application to AWS. The third party supplies a supported Docker image from a public registry. The image can run in as many containers as required to generate the route map.

The company has divided the delivery area into sections with supply hubs so that delivery drivers travel the shortest distance possible from the hubs to the customers. To reduce the time necessary to generate route maps, each section uses its own set of Docker containers with a custom configuration that processes orders only in the section's area.

The company needs the ability to allocate resources cost-effectively based on the number of running containers.

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

A. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2. Use the Amazon EKS CLI to launch the planning application in pods by using the --tags option to assign a custom tag to the pod.

B. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on AWS Fargate. Use the Amazon EKS CLI to launch the planning application. Use the AWS CLI tag-resource API call to assign a custom tag to the pod.

C. Create an Amazon Elastic Container Service (Amazon ECS) cluster on Amazon EC2. Use the AWS CLI with run-tasks set to true to launch the planning application by using the --tags option to assign a custom tag to the task.

D. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate. Use the AWS CLI run-task command and set enableECSManagedTags to true to launch the planning application. Use the --tags option to assign a custom tag to the task.

Suggested Answer: *D*

Community vote distribution

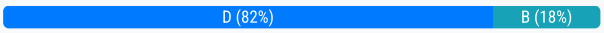

dev112233xx Highly Voted 2 years, 2 months ago

Selected Answer: D

D is the correct answer, When you use the APIs to create a service or run a task, you must set enableECSManagedTags to true for run-task and create-service. (see link below)
B doesn't make sense because EKS is more for complex orchestrated microservices apps, i don't think it needed in such scenario

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-using-tags.html

upvoted 16 times

Jay_2pt0_1 1 year, 7 months ago

Stepped through that same thought process

upvoted 1 times

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: D

EKS with Fargate is a more complex platform than ECS with Fargate. Kubernetes has a steeper learning curve than ECS, and requires more expertise to manage. ECS with Fargate is designed to be simple and easy to use, making it a good choice for organizations that want to quickly deploy containerized applications without having to manage the complexity of Kubernetes.

upvoted 6 times

0dc6cac Most Recent 2 weeks, 1 day ago

Selected Answer: D

it's D, people who pick B have never worked with EKS in a prod environment :-/

upvoted 1 times

altonh 5 months, 2 weeks ago

Selected Answer: D

For option B, how do you tag a POD using AWS CLI?

upvoted 1 times

amministrazione 10 months ago

D. Create an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate. Use the AWS CLI run-task command and set enableECSManagedTags to true to launch the planning application. Use the --tags option to assign a custom tag to the task.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: D

A and B = between EKS and ECS if K8s is not required I go for ECS

C = between EC2 and Fargate if nothing points you clearly to Ec2 i would go for Fargate (less overhead, could cost less)

D = correct

upvoted 3 times

ele 1 year, 5 months ago

Selected Answer: B

D is a trap, even if it's tempting, but '--tags ' is not a valid option for tagging ecs tasks/services.

B is the right answer.

upvoted 1 times

helloworldabc 10 months ago

just D

upvoted 1 times

cox1960 1 year, 5 months ago

--tags is valid

https://awscli.amazonaws.com/v2/documentation/api/latest/reference/ecs/run-task.html

upvoted 1 times

cox1960 1 year, 5 months ago

but --enable-ecs-managed-tags is the right option instead of "`enableECSManagedTags` to `true`"

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: D

D is best option with least operational overhead.

upvoted 3 times

atirado 1 year, 6 months ago

Selected Answer: D

Options A and C are more operationally complex than B and D because you will need to manage the EC2 instances and underpin the EKS cluster and the ECS service definition. And as if to make the selection easier, B and D explicitly mention using AWS Fargate in a way that works.

Selecting between Options B and D boils down the interpretation of “each section uses its own set of Docker containers with a custom configuration that processes orders only in the section's area”. The only indication in the question that kind of helps is "The third party supplies a supported Docker image from a public registry": The custom configuration is just for processing orders in the section's area rather something in the docker image itself.

upvoted 2 times

n_d1 1 year, 9 months ago

Selected Answer: D

As per the Amazon EKS documentation, the following EKS resources support tags:

- clusters
- managed node groups
- Fargate profiles

I think that rules out B in favour of D!

https://docs.aws.amazon.com/eks/latest/userguide/eks-using-tags.html#tag-resources

upvoted 3 times

Ganshank 1 year, 10 months ago

Real-world answer - B.

Certification answer - D.

upvoted 5 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

going with D

upvoted 2 times

rbm2023 2 years, 1 month ago

Selected Answer: D

Since the question where the requirement is the least operational overhead and we are between EKS and ECS, I would go for ECS, I believe EKS has more operational overhead for deploying and for operating. Also, you would probably have to apply less steps to build this structure using ECS when comparing with EKS.

upvoted 4 times

iamunstopable 2 years, 2 months ago

B is correct

Anytime you need Docker containers with a custom configuration use EKS

upvoted 2 times

Jay_2pt0_1 2 years, 2 months ago

Selected Answer: B

Like many have already stated, the debate is between B and D. I think B is the answer as "each section uses its own set of DOcker Containers with a customer configuration, " which leads me to believe that EKS orchestration is worthwhile in terms of operational overhead.

upvoted 3 times

mfsec 2 years, 3 months ago

Selected Answer: D

D is easier

upvoted 2 times

taer 2 years, 3 months ago

Selected Answer: D

I vote for D

upvoted 2 times

Question #128 *Topic 1*

A software company hosts an application on AWS with resources in multiple AWS accounts and Regions. The application runs on a group of Amazon EC2 instances in an application VPC located in the us-east-1 Region with an IPv4 CIDR block of 10.10.0.0/16. In a different AWS account, a shared services VPC is located in the us-east-2 Region with an IPv4 CIDR block of 10.10.10.0/24. When a cloud engineer uses AWS CloudFormation to attempt to peer the application VPC with the shared services VPC, an error message indicates a peering failure.

Which factors could cause this error? (Choose two.)

A. The IPv4 CIDR ranges of the two VPCs overlap

B. The VPCs are not in the same Region

C. One or both accounts do not have access to an Internet gateway

D. One of the VPCs was not shared through AWS Resource Access Manager

E. The IAM role in the peer accepter account does not have the correct permissions

Suggested Answer: *AE*

Community vote distribution

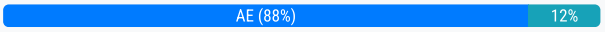

Appon Highly Voted 2 years, 4 months ago

Selected Answer: AE

https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-vpc-peering-error/

upvoted 9 times

zhangyu20000 Highly Voted 2 years, 5 months ago

AE is correct

D is not correct because you cannot share VPC via RAM, subnet can

upvoted 5 times

djeong95 1 year, 3 months ago

In this link, you can find VPC sharing being described as "In this model, the account that owns the VPC (owner) shares one or more subnets with other accounts (participants) that belong to the same organization from AWS Organization". You can share subnets using AWS RAM. I think it is safe to conclude you can share VPCs using RAM.

https://docs.aws.amazon.com/vpc/latest/userguide/vpc-sharing.html#vpc-share-prerequisites

upvoted 1 times

Syre Most Recent 9 months ago

Selected Answer: AB

E is wrong

upvoted 3 times

amministrazione 10 months ago

A. The IPv4 CIDR ranges of the two VPCs overlap

E. The IAM role in the peer accepter account does not have the correct permissions

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: AE

Option A and E

upvoted 1 times

m1xa 1 year, 7 months ago

Selected Answer: AE

https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html

https://repost.aws/knowledge-center/cloudformation-vpc-peering-error

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: AE

This is correct, per Appon's link
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: AE

AE for sure
upvoted 1 times

ThaiNT 2 years, 1 month ago

Selected Answer: BE

VPCs are not in the same Region.
upvoted 3 times

ThaiNT 2 years, 1 month ago

My bad, option B is incorrect.
upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: AE

AE is the best choice
upvoted 2 times

God_Is_Love 2 years, 3 months ago

Selected Answer: AE

FYI, Other reasons for issue :
If the IAM role in the accepter account doesn't have the right permissions

If the PeerRoleArn property isn't passed correctly when you create a VPC peering connection between VPCs in different accounts

If the PeerRegion property isn't passed correctly when you're creating a VPC peering connection between VPCs in different AWS Regions
upvoted 4 times

zozza2023 2 years, 5 months ago

Selected Answer: AE

A and E
upvoted 1 times

masetromain 2 years, 5 months ago

Selected Answer: AE

A is correct because the IPv4 CIDR ranges of the two VPCs overlap. The two VPCs have an IP range of 10.10.0.0/16 and 10.10.10.0/24, which means that they share the same 10.10.0.0 network. This causes a conflict in routing and will prevent the VPCs from being able to communicate with each other.

E is correct because the IAM role in the peer accepter account does not have the correct permissions. The role must have permissions to create, modify, and delete VPC peering connections in order for the peering to be established.

B, C, and D are not correct. The VPCs are in the same region, both accounts have access to an internet gateway and both VPCs are not shared through AWS Resource Access Manager.
upvoted 3 times

clownfishman 2 years ago

us-east-1 is in virginia, us-east-2 is in ohio - they are separate regions
upvoted 5 times

Arnaud92 1 year, 9 months ago

stop asking to ChatGPT
upvoted 7 times

m1xa 1 year, 7 months ago

It doesn't matter if both accounts are in the same region or not.
>>> The VPCs can be in different Regions (also known as an inter-Region VPC peering connection).
https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html
upvoted 1 times

Question #129 *Topic 1*

An external audit of a company's serverless application reveals IAM policies that grant too many permissions. These policies are attached to the company's AWS Lambda execution roles. Hundreds of the company's Lambda functions have broad access permissions such as full access to Amazon S3 buckets and Amazon DynamoDB tables. The company wants each function to have only the minimum permissions that the function needs to complete its task.

A solutions architect must determine which permissions each Lambda function needs.

What should the solutions architect do to meet this requirement with the LEAST amount of effort?

A. Set up Amazon CodeGuru to profile the Lambda functions and search for AWS API calls. Create an inventory of the required API calls and resources for each Lambda function. Create new IAM access policies for each Lambda function. Review the new policies to ensure that they meet the company's business requirements.

B. Turn on AWS CloudTrail logging for the AWS account. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail log. Review the generated policies to ensure that they meet the company's business requirements.

C. Turn on AWS CloudTrail logging for the AWS account. Create a script to parse the CloudTrail log, search for AWS API calls by Lambda execution role, and create a summary report. Review the report. Create IAM access policies that provide more restrictive permissions for each Lambda function.

D. Turn on AWS CloudTrail logging for the AWS account. Export the CloudTrail logs to Amazon S3. Use Amazon EMR to process the CloudTrail logs in Amazon S3 and produce a report of API calls and resources used by each execution role. Create a new IAM access policy for each role. Export the generated roles to an S3 bucket. Review the generated policies to ensure that they meet the company's business requirements.

Suggested Answer: *B*

Community vote distribution

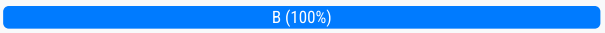

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: B

Access Analyzer uses automated reasoning to analyze resource policies and detect issues such as overly permissive access or violations of organizational security policies. It works by examining the policies attached to AWS resources, such as S3 buckets, IAM roles, and KMS keys, and identifying any potential security risks or policy violations.

upvoted 14 times

God_Is_Love 2 years, 3 months ago

fyi

ML tool - CodeGuru has two main components: CodeGuru Reviewer and CodeGuru Profiler.

CodeGuru Reviewer is a code review service that uses machine learning to identify code quality issues and security vulnerabilities in your application's source code. It analyzes the code and provides recommendations for improvements based on best practices, industry standards, and AWS experience.

CodeGuru Profiler is a profiling tool that uses machine learning to identify performance issues in your application code at runtime. It continuously analyzes the performance characteristics of your application code and provides recommendations for optimization.

upvoted 7 times

amministrazione Most Recent 10 months ago

B. Turn on AWS CloudTrail logging for the AWS account. Use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail log. Review the generated policies to ensure that they meet the company's business requirements.

upvoted 1 times

cox1960 1 year, 5 months ago

poor since B only works when functions are actually triggered and all the branches of the code are covered.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: B

Option B is obvious choice

upvoted 1 times

atirado 1 year, 6 months ago

Selected Answer: B

When approaching questions related to access permissions, it will always help to determine who is accessing what, in this case, it is Lambda functions accessing AWS services (S3 buckets and DynamoDB table).

The choice between A,B and C,D is then based on knowing that Code Guru and Access Analyzer used an automated process to detect issues in code and to compare actual access versus permissions - least effort than C & D.

That last bit is where the kicker is. The question refers to IAM execution roles with too-broad AWS IAM permissions to access AWS services and resources: You are looking for the option that tightens IAM policies rather than in AWS Lambda Function code.

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

B - basic access analyzer use case

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

keyword == Access Management Access Analyzer to generate IAM

upvoted 1 times

Alabi 2 years ago

Selected Answer: B

B definitely

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: B

B - Identity and Access Management Access Analyzer

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: B

Identity and Access Management Access Analyzer

upvoted 1 times

masetromain 2 years, 5 months ago

Selected Answer: B

The correct answer is B. Turn on AWS CloudTrail logging for the AWS account, and use AWS Identity and Access Management Access Analyzer to generate IAM access policies based on the activity recorded in the CloudTrail log. Review the generated policies to ensure that they meet the company's business requirements.

This is the least amount of effort as it makes use of AWS services that can automatically analyze the CloudTrail logs, generate the IAM policies, and provide a report for the review process.
Option A and D both involve additional steps such as running scripts or using Amazon EMR, which would take more effort to set up and maintain.
Option C is similar to option A and D but doesn't use any AWS services to help with the process.

upvoted 3 times

zhangyu20000 2 years, 5 months ago

B is correct

upvoted 1 times

Question #130 *Topic 1*

A solutions architect must analyze a company's Amazon EC2 instances and Amazon Elastic Block Store (Amazon EBS) volumes to determine whether the company is using resources efficiently. The company is running several large, high-memory EC2 instances to host database clusters that are deployed in active/passive configurations. The utilization of these EC2 instances varies by the applications that use the databases, and the company has not identified a pattern.

The solutions architect must analyze the environment and take action based on the findings.

Which solution meets these requirements MOST cost-effectively?

A. Create a dashboard by using AWS Systems Manager OpsCenter. Configure visualizations for Amazon CloudWatch metrics that are associated with the EC2 instances and their EBS volumes. Review the dashboard periodically, and identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.

B. Turn on Amazon CloudWatch detailed monitoring for the EC2 instances and their EBS volumes. Create and review a dashboard that is based on the metrics. Identify usage patterns. Rightsize the EC2 instances based on the peaks in the metrics.

C. Install the Amazon CloudWatch agent on each of the EC2 instances. Turn on AWS Compute Optimizer, and let it run for at least 12 hours. Review the recommendations from Compute Optimizer, and rightsize the EC2 instances as directed.

D. Sign up for the AWS Enterprise Support plan. Turn on AWS Trusted Advisor. Wait 12 hours. Review the recommendations from Trusted Advisor, and rightsize the EC2 instances as directed.

Suggested Answer: *C*

Community vote distribution

C (97%)

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: C

AWS Compute Optimize helps analyze the usage patterns of AWS resources, such as EC2 instances and Auto Scaling groups, and makes recommendations on how to optimize them for performance and cost using machine learning algorithms.It then generates recommendations that can be used to adjust instance types, purchase options, and other parameters.It provides two types of recommendations:
Recommended instance types - recommends instance types that are more cost-effective and better suited to the workload requirements.
Recommended purchase options - recommends purchasing options, such as Reserved Instances or Savings Plans, that can help customers save money on their compute resources.

upvoted 17 times

God_Is_Love 2 years, 3 months ago

A is wrong.
OpsCenter, a capability of AWS Systems Manager, provides a central location where operations engineers and IT professionals can manage operational work items (OpsItems) related to AWS resources. An OpsItem is any operational issue or interruption that needs investigation and remediation. Using OpsCenter, you can view contextual investigation data about each OpsItem, including related OpsItems and related resources. You can also run Systems Manager Automation runbooks to resolve OpsItems.

upvoted 4 times

God_Is_Love 2 years, 3 months ago

fyi Pricing looks cheap too - https://aws.amazon.com/compute-optimizer/pricing/

upvoted 2 times

amministrazione Most Recent 10 months ago

C. Install the Amazon CloudWatch agent on each of the EC2 instances. Turn on AWS Compute Optimizer, and let it run for at least 12 hours. Review the recommendations from Compute Optimizer, and rightsize the EC2 instances as directed.

upvoted 1 times

saggy4 1 year, 4 months ago

Selected Answer: C

A - Not possible
D - Costliest Option possible
now between B and C

The question mentions high-memory EC2 instances.
You cannot get memory metrics without the Cloudwatch agent installed hence C.
upvoted 2 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C is most cost effective choice.
upvoted 1 times

wmp7039 1 year, 6 months ago

C is incorrect : When you first opt in Compute Optimizer, it may take up to 24 hours to fully analyze the AWS resources in your account.
https://aws.amazon.com/compute-optimizer/faqs/
upvoted 1 times

carpa_jo 1 year, 6 months ago

You are correct that in the FAQ you've linked it says 24 hours, but in other places of the AWS documentation it says 12 hours, like here:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/asg-getting-recommendations.html#viewing-recommendations
or here: https://docs.aws.amazon.com/awssupport/latest/user/compute-optimizer-with-trusted-advisor.html
Seems like even AWS doesn't know :D So I would still go with C.
upvoted 1 times

atirado 1 year, 6 months ago

Selected Answer: C

Option A is not in the running because it will require incurring further expense to address the cost issue.

Option D is expensive - the Enterprise Support plan charges a minimum flat fee minimum or a % of your AWS bill. This could be a large amount for the company's hundreds of instances.

Option B is expensive - Detailed monitoring scales based on the number of metrics and the number of resources. The company has hundreds of instances so this option could potentially be more expensive than D.

Option C - Compute Optimizer will provide improvement suggestions based on 14 prior days usage data from the moment it was enabled. Moreover, the default service option is free. Nothing is said about the custom metrics being used for the CloudWatch agent but it could be the most expensive of all options if mis-used. So either cost 0 or incredibly large if used carelessly.
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C. need CW agent for RAm util
upvoted 1 times

Fredonly 2 years, 2 months ago

Selected Answer: C

C- Compute Optimizer is the easiest solution
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

C - cost optimizer
upvoted 1 times

mfsec 2 years, 3 months ago

*Compute
upvoted 1 times

spd 2 years, 3 months ago

Selected Answer: C

C is correct - Optimzer
upvoted 2 times

kiran15789 2 years, 4 months ago

Selected Answer: A

Option C may be a good solution to rightsize the EC2 instances but may incur additional cost for installing the Amazon CloudWatch agent on each of the EC2 instances.

The MOST cost-effective solution to analyze the company's Amazon EC2 instances and Amazon EBS volumes is to create a dashboard using AWS Systems Manager OpsCenter. The OpsCenter dashboard can be configured to visualize the Amazon CloudWatch metrics associated with the EC2 instances and their EBS volumes. By reviewing the dashboard periodically, usage patterns can be identified, and EC2 instances can be right-sized based on the peaks in the metrics.

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Bro, install cost is 0. Simple linux command > sudo yum install amazon-cloudwatch-agent

upvoted 2 times

masetromain 2 years, 5 months ago

Selected Answer: C

The correct answer is C. Installing the Amazon CloudWatch agent on each of the EC2 instances and turning on AWS Compute Optimizer allows the solutions architect to analyze the environment and make recommendations on the sizing of the EC2 instances in a cost-effective way. AWS Compute Optimizer analyzes the utilization of the instances and recommends the optimal instance types for the workloads. This solution is more cost-effective than creating a dashboard and reviewing it periodically, or signing up for the AWS Enterprise Support plan and waiting for Trusted Advisor recommendations.

upvoted 3 times

zhangyu20000 2 years, 5 months ago

C is correct, with computer optimizer

upvoted 1 times

Question #131 *Topic 1*

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company uses AWS Control Tower for governance and uses AWS Transit Gateway for VPC connectivity across accounts.

In an AWS application account, the company's application team has deployed a web application that uses AWS Lambda and Amazon RDS. The company's database administrators have a separate DBA account and use the account to centrally manage all the databases across the organization. The database administrators use an Amazon EC2 instance that is deployed in the DBA account to access an RDS database that is deployed m the application account.

The application team has stored the database credentials as secrets in AWS Secrets Manager in the application account. The application team is manually sharing the secrets with the database administrators. The secrets are encrypted by the default AWS managed key for Secrets Manager in the application account. A solutions architect needs to implement a solution that gives the database administrators access to the database and eliminates the need to manually share the secrets.

Which solution will meet these requirements?

A. Use AWS Resource Access Manager (AWS RAM) to share the secrets from the application account with the DBA account. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the shared secrets. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

B. In the application account, create an IAM role that is named DBA-Secret. Grant the role the required permissions to access the secrets. In the DBA account, create an IAM role that is named DBA-Admin. Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets

C. In the DBA account create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the secrets and the default AWS managed key in the application account. In the application account, attach resource-based policies to the key to allow access from the DBA account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

D. In the DBA account, create an IAM role that is named DBA-Admin. Grant the role the required permissions to access the secrets in the application account. Attach an SCP to the application account to allow access to the secrets from the DBA account. Attach the DBA-Admin role to the EC2 instance for access to the cross-account secrets.

Suggested Answer: *A*

Community vote distribution

bititan Highly Voted 1 year, 10 months ago

Selected Answer: B

Follow below link. It has both option to be used for this scenarios. But default kms key can not be used so B
https://aws.amazon.com/blogs/database/design-patterns-to-access-cross-account-secrets-stored-in-aws-secrets-manager/
upvoted 15 times

Sarutobi Highly Voted 1 year, 8 months ago

Selected Answer: B

Although I think B is the best, it is missing to mention of the trust policy in the application account.
upvoted 6 times

ninomfr64 11 months, 1 week ago

Grant the DBA-Admin role the required permissions to assume the DBA-Secret role in the application account. This sounds like a trust policy to me
upvoted 1 times

ninomfr64 Most Recent 11 months, 1 week ago

Selected Answer: B

A = Secret is not a RAM sharable resource. But who can recall this full list? Thus my reasoning is, I would expect more details for sharing via RAM like enable AWS Org sharing, assign permission (actions allowed on the shared resource) and select the external principal.
B = correct see https://aws.amazon.com/blogs/database/design-patterns-to-access-cross-account-secrets-stored-in-aws-secrets-manager/
C = cannot cross-account access AWS managed KMS key as you do not have control on key policy
D = SCP can only remove permissions. Even tough an SCP doesn't prevent you from accessing a secret, you still need to have IAM user permission and/or resource based policy in place to actually access

upvoted 4 times

horyoryo 1 year ago

option b

upvoted 1 times

career360guru 1 year ago

Selected Answer: B

Option B

upvoted 1 times

bjexamprep 1 year ago

Selected Answer: B

Even B is the best answer among all the options, actually B is not correct. Without permission to access the KMS key, B cannot decrypt the secret.

upvoted 2 times

bjexamprep 9 months, 3 weeks ago

I was wrong. It is using AWS managed default encryption key, so it doesn't need the permission to access KMS key. The flaw of B is trust relationship policy.

upvoted 1 times

severlight 1 year, 1 month ago

Selected Answer: B

the Secrets Manager keys cannot be shared with RAM, key policy(resource policy) for the default KMS key managed by AWS cannot be changed, role is identity and can be granted access to assume other role

upvoted 1 times

rlf 1 year, 2 months ago

Answer is B.

Option A is wrong. AWS RAM can not share AWS Secrets Manager (see shareable resources in https://docs.aws.amazon.com/ram/latest/userguide/shareable.html)

upvoted 3 times

uC6rW1aB 1 year, 3 months ago

Selected Answer: A

Both Option A and Option B give repository administrators access to the repository and eliminate the need to manually share secrets.

Option A is a relatively simple process of sharing secrets with AWS RAM and setting up an IAM role within the DBA account.

Option B requires creating an IAM role in two different AWS accounts and setting cross-account permissions, which is a more complicated process.

So, while both A and B accomplish the goal, option A is simpler and more straightforward.

upvoted 1 times

chikorita 1 year, 3 months ago

who said we can share secrets using RAM??

i just checked under RAM and allowed sharable AWS services

AWS Secrets Manager is NOT one of those

Answer is B

upvoted 4 times

venvig 1 year, 4 months ago

Selected Answer: B

As several people have highlighted, we refer to the blog https://aws.amazon.com/blogs/database/design-patterns-to-access-cross-account-secrets-stored-in-aws-secrets-manager/

Want to provide the following comment to emphasize why "C" is NOT even possible.

In Option C, its mentioned that the default AWS Managed CMK is used by the secrets manager.

We cannot provide any custom permissions to the AWS Managed CMK and by extension, its not possible to allow cross account access to it. So, only Option B is valid.

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

its a b

upvoted 1 times

Jackhemo 1 year, 6 months ago

Guys, you want to know the right answer? Copy paste the whole question to olabiba.ai
The answer is B
upvoted 1 times

OCHT 1 year, 8 months ago

Selected Answer: A

Option A is the correct answer because it meets the requirement of giving the database administrators access to the database and eliminates the need to manually share the secrets. AWS Resource Access Manager (AWS RAM) enables you to share AWS resources with other accounts within your organization or organizational units (OUs) in AWS Organizations. By using AWS RAM to share the secrets from the application account with the DBA account, you can eliminate the need for manual sharing of secrets.

Option B involves creating an IAM role in the application account and another IAM role in the DBA account. The DBA-Admin role in the DBA account would need to assume the DBA-Secret role in the application account to access the secrets. This approach adds complexity and does not eliminate the need for manual sharing of secrets.

In summary, Option A is a simpler and more efficient solution that meets the requirements.
upvoted 2 times

Maria2023 1 year, 6 months ago

I couldn't find any option to share Secret Manager resources via RAM, did anyone try it?
upvoted 4 times

dev112233xx 1 year, 8 months ago

Selected Answer: B

B is correct, D doesn't make sense! SCP doesn't give any permission.. it just defines what can be allowed. you still need an IAM role/policy
upvoted 2 times

mfsec 1 year, 9 months ago

Selected Answer: B

B is the best choice
upvoted 2 times

DWsk 1 year, 9 months ago

Selected Answer: B

Has to be B because C is not possible.
I get that you can't share access to the default KMS key, but how does it work to share access through a cross account role? How does the role in the DBA account decrypt the secrets that are encrypted by the default key if the role doesn't have permissions to that key?
upvoted 4 times

kiran15789 1 year, 9 months ago

Selected Answer: B

cross account assume role
upvoted 2 times

Question #132 *Topic 1*

A company manages multiple AWS accounts by using AWS Organizations. Under the root OU, the company has two OUs: Research and DataOps.

Because of regulatory requirements, all resources that the company deploys in the organization must reside in the ap-northeast-1 Region. Additionally, EC2 instances that the company deploys in the DataOps OU must use a predefined list of instance types.

A solutions architect must implement a solution that applies these restrictions. The solution must maximize operational efficiency and must minimize ongoing maintenance.

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

A. Create an IAM role in one account under the DataOps OU. Use the ec2:InstanceType condition key in an inline policy on the role to restrict access to specific instance type.

B. Create an IAM user in all accounts under the root OU. Use the aws:RequestedRegion condition key in an inline policy on each user to restrict access to all AWS Regions except ap-northeast-1.

C. Create an SCP. Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root OU.

D. Create an SCP. Use the ec2:Region condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root OU, the DataOps OU, and the Research OU.

E. Create an SCP. Use the ec2:InstanceType condition key to restrict access to specific instance types. Apply the SCP to the DataOps OU.

Suggested Answer: *CE*

Community vote distribution

CE (100%)

OCHT Highly Voted 1 year, 8 months ago

Selected Answer: CE

C. Create an SCP. Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root OU. This will ensure that all resources deployed in the organization reside in the ap-northeast-1 Region.

E. Create an SCP. Use the ec2:InstanceType condition key to restrict access to specific instance types. Apply the SCP to the DataOps OU. This will ensure that EC2 instances deployed in the DataOps OU use only the predefined list of instance types.

upvoted 5 times

OCHT 1 year, 8 months ago

Option D is incorrect because it suggests using the ec2:Region condition key to restrict access to all AWS Regions except ap-northeast-1. However, the ec2:Region condition key is not a valid condition key for EC2 actions. Instead, the aws:RequestedRegion condition key should be used to restrict access to specific AWS Regions.

Additionally, applying the SCP to the root OU, the DataOps OU, and the Research OU is unnecessary because applying the SCP to the root OU alone will ensure that the restriction applies to all accounts in the organization, including those in the DataOps and Research OUs.

In summary, option D is incorrect because it suggests using an invalid condition key and because applying the SCP to multiple OUs is unnecessary.

upvoted 3 times

career360guru Most Recent 1 year ago

Selected Answer: CE

Option C & E

upvoted 1 times

venvig 1 year, 4 months ago

Selected Answer: CE

Very straightforward

upvoted 2 times

dtha1002 1 year, 5 months ago

Selected Answer: CE

C for all resources region

and E for DataOps OU launch instantce type

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: CE

its CE

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: CE

SCP's are the most efficient here

upvoted 1 times

tatdatpham 1 year, 11 months ago

Selected Answer: CE

With AWS Org, consider SCP first.

In this scenario, Only C,D,E are mention about SCP, but D apply for all, not only the DataOps OU

upvoted 4 times

masetromain 1 year, 11 months ago

Selected Answer: CE

The correct options are C and E.

Option C: Create an SCP. Use the aws:RequestedRegion condition key to restrict access to all AWS Regions except ap-northeast-1. Apply the SCP to the root OU.

This option is correct because it allows the company to restrict access to all AWS regions except for ap-northeast-1. This ensures that all resources deployed in the organization must reside in the ap-northeast-1 region. By applying the SCP to the root OU, it ensures that all accounts and OUs under the root will be affected.

Option E: Create an SCP. Use the ec2:InstanceType condition key to restrict access to specific instance types. Apply the SCP to the DataOps OU.

This option is correct because it allows the company to restrict access to specific instance types, which is required for the DataOps OU. By applying the SCP to the DataOps OU, it ensures that only resources deployed in the DataOps OU will be affected by the restriction.

upvoted 4 times

masetromain 1 year, 11 months ago

Option A is incorrect because it only restricts access to specific instance types, but it does not restrict access to a specific region.

Option B is incorrect because it is applied to IAM users rather than OUs, which would not effectively apply the restriction to all resources in the organization.

Option D is incorrect because it uses the ec2:Region condition key which would not allow to restrict the instances types only in the DataOps OU.

By creating an SCP that uses the aws:RequestedRegion condition key and restricting access to all regions except ap-northeast-1 and applying it to the root OU, this ensures that all resources deployed in the organization will reside in the ap-northeast-1 Region.

By creating an SCP that uses the ec2:InstanceType condition key and restricts access to specific instance types and applying it to the DataOps OU, this ensures that all EC2 instances deployed in the DataOps OU will use the predefined list of instance types.

upvoted 1 times

zhangyu20000 1 year, 11 months ago

CE is correct

upvoted 1 times

Question #133 *Topic 1*

A company runs a serverless application in a single AWS Region. The application accesses external URLs and extracts metadata from those sites. The company uses an Amazon Simple Notification Service (Amazon SNS) topic to publish URLs to an Amazon Simple Queue Service (Amazon SQS) queue. An AWS Lambda function uses the queue as an event source and processes the URLs from the queue. Results are saved to an Amazon S3 bucket.

The company wants to process each URL in other Regions to compare possible differences in site localization. URLs must be published from the existing Region. Results must be written to the existing S3 bucket in the current Region.

Which combination of changes will produce multi-Region deployment that meets these requirements? (Choose two.)

A. Deploy the SQS queue with the Lambda function to other Regions.

B. Subscribe the SNS topic in each Region to the SQS queue.

C. Subscribe the SQS queue in each Region to the SNS topic.

D. Configure the SQS queue to publish URLs to SNS topics in each Region.

E. Deploy the SNS topic and the Lambda function to other Regions.

Suggested Answer: *AC*

Community vote distribution

AC (100%)

SK_Tyagi Highly Voted 1 year, 4 months ago

Selected Answer: AC

SNS being the publisher, SQS is subscribing

upvoted 6 times

rlf Highly Voted 1 year, 2 months ago

AC.

Amazon SNS supports cross-region deliveries.

https://docs.aws.amazon.com/sns/latest/dg/sns-cross-region-delivery.html

upvoted 5 times

SeemaDataReader Most Recent 11 months, 1 week ago

Selected Answer: AC

SNS in Region A, SQS + Lambda in Region A & B, S3 Bucket in Region A

upvoted 2 times

career360guru 1 year ago

Selected Answer: AC

A and C

upvoted 1 times

Passexam4sure_com 1 year, 2 months ago

Selected Answer: AC

Deploy the SQS queue with the Lambda function to other Regions.

Subscribe the SQS queue in each Region to the SNS topic.

upvoted 3 times

NikkyDicky 1 year, 5 months ago

Selected Answer: AC

It's an AC

upvoted 2 times

Maria2023 1 year, 6 months ago

Selected Answer: AC

Basically, you need to replicate it all except the bucket in the other regions. The question is explained very vaguely however

upvoted 3 times

awsleffe 8 months, 3 weeks ago

SNS is the publisher and must stay in same region

upvoted 1 times

Parsons 1 year, 8 months ago

Selected Answer: AC

A, C is correct.

It looks like Fan out pattern.

upvoted 3 times

Kampton 1 year, 8 months ago

Why would need to deploy SQS with Lambda? Makes no sense! It's BE.

upvoted 1 times

Diego1414 1 year, 7 months ago

It's SNS that publishes not SQS

upvoted 2 times

Asagumo 1 year, 8 months ago

What does it mean in Option A that Lambda deploys SQS?

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: AC

AC - SQS

upvoted 2 times

Zek 1 year, 9 months ago

support A,C. https://www.examtopics.com/discussions/amazon/view/74009-exam-aws-certified-solutions-architect-professional-topic-1/

upvoted 1 times

MasterP007 1 year, 10 months ago

A & C - Deploy & Subscribe SQS.

upvoted 1 times

zozza2023 1 year, 11 months ago

Selected Answer: AC

A and C

upvoted 3 times

masetromain 1 year, 11 months ago

Selected Answer: AC

Option A is correct because deploying the SQS queue with the Lambda function to other regions will allow the application to process URLs in those regions and compare differences in site localization.

Option C is correct because subscribing the SQS queue in each region to the SNS topic in the existing region will allow the application to publish URLs to the existing SNS topic and have those URLs processed in other regions.

Option B is incorrect because subscribing the SNS topic in each region to the SQS queue in the existing region would not allow URLs to be processed in other regions.

Option D is incorrect because configuring the SQS queue to publish URLs to SNS topics in each region would not ensure that the URLs are processed in those regions.

Option E is incorrect because deploying the SNS topic and Lambda function to other regions without the SQS queue would not allow the application to process URLs in those regions.

upvoted 4 times

zhangyu20000 1 year, 11 months ago

AC is correct

upvoted 1 times

Question #134 *Topic 1*

A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instances. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM, and is highly CPU intensive. The application is scheduled to run every 4 hours and runs for up to 20 minutes. A solutions architect wants to revise the architecture for the solution.

Which strategy should the solutions architect use?

A. Use AWS Lambda to run the application. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours.

B. Use AWS Batch to run the application. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.

C. Use AWS Fargate to run the application. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours.

D. Use Amazon EC2 Spot Instances to run the application. Use AWS CodeDeploy to deploy and run the application every 4 hours.

Suggested Answer: *C*

Community vote distribution

C (89%) 11%

zhangyu20000 Highly Voted 1 year, 11 months ago

C is correct. only eventbridge can run scheduled task

upvoted 16 times

Maria2023 Highly Voted 1 year, 6 months ago

Selected Answer: C

If there wasn't a schedule element I would choose AWS Batch because it pretty much loads a container and does the job, especially since it's like a 20-minute job. However the step functions part doesn't help with the scheduling part, hence I go for C

upvoted 6 times

eesa Most Recent 2 months, 1 week ago

Selected Answer: C

AWS Fargate is ideal for running containerized workloads without managing underlying EC2 instances. Even though the application is a binary and its source code cannot be modified, it can be easily packaged into a Docker container without changing the binary itself.

Since the application runs periodically (every 4 hours) and for a short duration (up to 20 minutes), Fargate provides a cost-effective, serverless execution environment.

Amazon EventBridge (CloudWatch Events) can be scheduled to invoke Fargate tasks precisely at defined intervals.

upvoted 1 times

Longc 3 months, 3 weeks ago

Selected Answer: C

Option C clearly includes EventBridge for scheduling, aligning with the requirement to run tasks every 4 hours. While AWS Batch is technically better for CPU-intensive workloads, the lack of explicit EventBridge integration in Option B makes C the correct answer under AWS's service design principles.

upvoted 1 times

albert_kuo 3 months, 3 weeks ago

Selected Answer: B

B. Use AWS Batch to run the application. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.

upvoted 2 times

9d7a975 4 months, 2 weeks ago

Selected Answer: B

B: Usa uma máquina de estado do AWS Step Functions para invocar o trabalho do AWS Batch a cada 4 horas.
Por que não é a letra C : Embora possa executar containers, é mais adequado para aplicações de longa duração

upvoted 3 times

TonytheTiger 8 months, 3 weeks ago

Option C : https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/run-event-driven-and-scheduled-workloads-at-scale-with-aws-fargate.html

upvoted 1 times

ninomfr64 11 months, 1 week ago

A = CW Log cannot invoke lambda every 4 hours

B = Step Function cannot invoke batch job every 4 hour (unless you use an EventBridhe scheduled event)

C = correct (but I do not like when Fargate is mentioned as a standalone service, as it is a serverless compute option for some some services)

D = CodeDeploy cannot run an application every 4 hours

upvoted 2 times

cox1960 11 months, 2 weeks ago

none. "highly CPU intensive" means no Fargate. scheduling means eventbridge.

upvoted 2 times

holymancolin 11 months, 3 weeks ago

https://aws.amazon.com/about-aws/whats-new/2018/10/aws-lambda-supports-functions-that-can-run-up-to-15-minutes/

Lambda's max running time is 15 mins, cannot support up to 20mins application.

upvoted 1 times

haha001 1 year ago

https://aws.amazon.com/tutorials/scheduling-a-serverless-workflow-step-functions-amazon-eventbridge-scheduler/

Step Function cannot schedule a job. Step Function needs EventBridge as the scheduler.

upvoted 2 times

career360guru 1 year ago

Selected Answer: C

B is not possible as Step Function can not be used to run scheduled a job every 4 hour

upvoted 1 times

task_7 1 year, 3 months ago

Selected Answer: D

containers are well-suited for applications that are built in microservices architecture, where each service is a self-contained unit that performs a specific task. These types of applications are typically designed to be scalable and easy to deploy, making them a good fit for containerization.

I feel D is the best option

upvoted 2 times

teo2157 10 months ago

you can´t garantee with spot instances that they're available every 4 hours, C is the answer

upvoted 3 times

uC6rW1aB 1 year, 3 months ago

Selected Answer: C

I think Both B C is missing some key point

Option B does not explain how to AWS Step Functions to trigger an AWS Batch job regually, in this case 4 hours per run.

Option C does not explain how to use EventBridge to call the Fargate task, which is not native support, it might involved lambda to achive.

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

C. schedule -> eventbridge

upvoted 1 times

rbm2023 1 year, 7 months ago

Selected Answer: C

The application is a Linux binary which can be packaged into a container, then run on AWS Fargate and scheduled using Event Bridge.

```
# Use a base image that matches your application's runtime environment
FROM ubuntu:latest
# Copy the Linux binary into the container
COPY myapp /usr/local/bin/myapp
# Set the entry point to execute the binary
ENTRYPOINT ["/usr/local/bin/myapp"]
```

upvoted 3 times

mfsec 1 year, 9 months ago

Selected Answer: C

C - Fargate is the best choice here

upvoted 1 times

Question #135 *Topic 1*

A company is creating a sequel for a popular online game. A large number of users from all over the world will play the game within the first week after launch. Currently, the game consists of the following components deployed in a single AWS Region:

• Amazon S3 bucket that stores game assets
• Amazon DynamoDB table that stores player scores

A solutions architect needs to design a multi-Region solution that will reduce latency, improve reliability, and require the least effort to implement.

What should the solutions architect do to meet these requirements?

A. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Cross-Region Replication. Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.

B. Create an Amazon CloudFront distribution to serve assets from the S3 bucket. Configure S3 Same-Region Replication. Create a new DynamoDB table in a new Region. Configure asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC).

C. Create another S3 bucket in a new Region, and configure S3 Cross-Region Replication between the buckets. Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region. Configure DynamoDB global tables by enabling Amazon DynamoDB Streams, and add a replica table in a new Region.

D. Create another S3 bucket in the sine Region, and configure S3 Same-Region Replication between the buckets. Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets. Create a new DynamoDB table in a new Region. Use the new table as a replica target for DynamoDB global tables.

Suggested Answer: *C*

Community vote distribution

zozza2023 Highly Voted 2 years, 5 months ago

Selected Answer: C

DynamoDB global tables + S3 replication+Cloudfront

upvoted 14 times

masetromain Highly Voted 2 years, 5 months ago

Option C is the correct answer because it meets the requirements of reducing latency, improving reliability and requiring minimal effort to implement.

By creating another S3 bucket in a new Region, and configuring S3 Cross-Region Replication between the buckets, the game assets will be replicated to the new Region, reducing latency for users accessing the assets from that region. Additionally, by creating an Amazon CloudFront distribution and configuring origin failover with two origins accessing the S3 buckets in each Region, it ensures that the game assets will be served to users even if one of the regions becomes unavailable.

Configuring DynamoDB global tables by enabling Amazon DynamoDB Streams, and adding a replica table in a new Region, will also improve reliability by allowing the player scores to be replicated and updated in multiple regions, ensuring that the scores are available even in the event of a regional failure.

upvoted 7 times

masetromain 2 years, 5 months ago

Option A is not correct because using the new table as a replica target for DynamoDB global tables will not improve reliability. The same applies for Option D, which only uses S3 Same-Region Replication, which will not reduce latency for users in other regions.

Option B is not correct because configuring asynchronous replication between the DynamoDB tables by using AWS Database Migration Service (AWS DMS) with change data capture (CDC) is not the best solution for this use case. It would require additional configuration and management effort.

upvoted 3 times

jimee11 Most Recent 1 month, 2 weeks ago

Selected Answer: C

CloudFront supports two origins, and Streaming is required to enable Global tables.

upvoted 1 times

Daniel76 8 months, 3 weeks ago

Selected Answer: C

Just to add for DynamoDB, indeed you will need to create replica in the new region when creating global table, making it accessible in the new region nearer to the user.

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/V2globaltables.tutorial.html

upvoted 1 times

JoeTromundo 8 months, 3 weeks ago

Selected Answer: C

Option C is correct.

Just to clarify: AWS uses DynamoDB Streams to replicate DynamoDB Global Tables. Using the Console, it is enabled automatically. Using the CLI, you must enable it explicitly by using StreamEnabled=true.

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

A = "Configure S3 Cross-Region Replication" but doesn't create a new bucket in another region.

B = "Configure S3 Same-Region Replication" without creating a second bucket and this should be cross-region. AWS DMS with CDC is not a good fit here, global table is the right option here

C = correct

D = we need the new bucket in a different region

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C

upvoted 2 times

shaaam80 1 year, 6 months ago

Selected Answer: C

Answer C.

Regarding DynamoDB Streams -

Global tables use DynamoDB Streams to replicate data across different Regions. When you create a replica for a global table, a stream is created by default. Any changes to a replica are replicated to all the other replicas within the same global table within a second using DynamoDB Streams.

upvoted 2 times

blackgamer 1 year, 7 months ago

The answer is A. C added unnecessary complexities such as Amazon DynamoDB Streams and Origin Failover.

upvoted 1 times

helloworldabc 10 months ago

just C

upvoted 1 times

ninomfr64 1 year, 5 months ago

Option A doesn't mention creating a new bucket in a different region

upvoted 1 times

Jay_2pt0_1 1 year, 7 months ago

I initially thought it was C, but I was torn between A and C. You may be right.

upvoted 1 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: A

other option are incorrect.

B: Configure S3 Same-Region Replication.---> It's not meet multi-region requirement.

C: Create an Amazon CloudFront distribution and configure origin failover with two origins accessing the S3 buckets in each Region. ---> It's not support for this kinda failover

D: Create another S3 bucket in the same Region, and configure S3 Same-Region Replication between the buckets. ---> It's not meet multi-region requirement.

upvoted 2 times

ninomfr64 1 year, 5 months ago

C is correct, Origin Group allows failover see
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html

upvoted 2 times

dkcloudguru 1 year, 9 months ago

option c is the easiest way to do

upvoted 1 times

ProMax 1 year, 10 months ago

Selected Answer: A

Creating an Amazon CloudFront distribution will reduce latency for global users by serving assets from the closest edge location. S3 Cross-Region Replication will ensure that game assets are available in another region, improving reliability. Creating a new DynamoDB table in a new region and using it as a replica target for DynamoDB global tables will enable multi-region replication, improving reliability.

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: C

Option C has another differentiator - DynamoDBStreams that will assist in Reliability

upvoted 2 times

ggrodskiy 1 year, 11 months ago

Correct A.
CloudFront does not support origin failover with two origins accessing the S3 buckets in each Region. According to the AWS documentationhttps://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html, origin failover only works within the same Region, not across Regions. This means that you can only configure origin failover with two origins that are in the same Region as the CloudFront distribution. If you want to use origin failover with S3 buckets in different Regions, you need to create multiple CloudFront distributions, one for each Region, and configure them to use the same domain name with geolocation routinghttps://blog.ippon.tech/when-a-cloudfront-origin-must-fail-for-testing-high-availability/.

upvoted 1 times

venvig 1 year, 10 months ago

Referred to your AWS doc link. I don't see any condition that states that the origins in the origin group cannot be from two different regions. Can you provide the statement from the AWS doc that you are referring to please ?

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

weird question wording, but C fit more

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

Create another S3 bucket in a new Region, and configure S3 Cross-Region Replication between the buckets

upvoted 2 times

zhangyu20000 2 years, 5 months ago

C is correct. S3 cross replicate, CloudFront, Dynamodb global database and origin failover

upvoted 2 times

Question #136 *Topic 1*

A company has an on-premises website application that provides real estate information for potential renters and buyers. The website uses a Java backend and a NoSQL MongoDB database to store subscriber data.

The company needs to migrate the entire application to AWS with a similar structure. The application must be deployed for high availability, and the company cannot make changes to the application.

Which solution will meet these requirements?

A. Use an Amazon Aurora DB cluster as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

B. Use MongoDB on Amazon EC2 instances as the database for the subscriber data. Deploy EC2 instances in an Auto Scaling group in a single Availability Zone for the Java backend application.

C. Configure Amazon DocumentDB (with MongoDB compatibility) with appropriately sized instances in multiple Availability Zones as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

D. Configure Amazon DocumentDB (with MongoDB compatibility) in on-demand capacity mode in multiple Availability Zones as the database for the subscriber data. Deploy Amazon EC2 instances in an Auto Scaling group across multiple Availability Zones for the Java backend application.

Suggested Answer: *D*

Community vote distribution

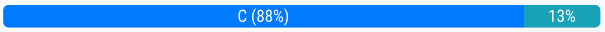

uC6rW1aB Highly Voted 1 year, 9 months ago

Selected Answer: C

C correct

DocumentDB only have on-demand instance but not on-demand capacity mode, the mode is for DynamoDB

upvoted 12 times

ninomfr64 Highly Voted 1 year, 5 months ago

Selected Answer: C

A = Aurora supports MySQL and PostgreSQL, not MongoDB. App changes are not allowed

B = This could work but DocumentDB provides managed MongoDB instance that is preferable

C = correct

D = there isn't on-demand capacity mode, in 2022 launched MondoDB Elastic Cluster that eliminates the need to choose, manage or upgrade instances and allows to scale up to 4PiB storage whereas instance based scales up to 128TiB.

I thing this question is pre elastic cluster as this is ambiguous between C and D

upvoted 5 times

jimee11 Most Recent 1 month, 2 weeks ago

Selected Answer: C

C: DocumentDB has an on-demand instance type but NO on-demand capacity mode. Note the difference between the two:

On-demand instance type is specific to EC2 pricing for hourly or per-second compute capacity. On-demand capacity mode is specific for DynamoDB and Kinese for pay-per-request pricing and no up-front capacity planning.

upvoted 1 times

cnethers 12 months ago

D is the correct answer https://aws.amazon.com/documentdb/pricing/

on-demand instance is supported by DocumentDB

upvoted 1 times

helloworldabc 10 months ago

just C

upvoted 2 times

gofavad926 1 year, 3 months ago

Selected Answer: C

C, documented. No exists the on-demand capacity mode

upvoted 1 times

AimarLeo 1 year, 5 months ago

'Appropriately sized instances' Means on-demand ? that is quite vague..

upvoted 3 times

jpa8300 1 year, 5 months ago

Selected Answer: D

DocumentDB does indeed support on-demand capacity mode (Contrary to what other users say here)
https://aws.amazon.com/blogs/database/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/
On-Demand is ideally to a use case where you have unpredictable or variable database workloads, like this case, it is not said anywhere the expected workload, so it is better to start with On-demand , and later when you know the workload you can cahnge it.

upvoted 2 times

buriz 1 year, 5 months ago

what you have linked here is a dynamodb article not a documentDB one, documentDB does not support on-demand capacity mode - https://aws.amazon.com/documentdb/faqs/

"You can scale the compute resources allocated to your instance in the AWS Management Console by selecting the desired instance and clicking the “modify” button. Memory and CPU resources are modified by changing your instance class."

upvoted 1 times

ninomfr64 1 year, 5 months ago

There is no on-demand capacity for DocumentDB, however Elastic Cluster option is provided "Elastic Clusters enables you to elastically scale your document database to handle millions of writes and reads, with petabytes of storage capacity" see https://aws.amazon.com/documentdb/faqs/#:~:text=to%20learn%20more.-,Elastic%20Clusters,-What%20is%20Amazon

upvoted 1 times

chicagobeef 1 year, 5 months ago

This is DynamoDB, not DocumentDB. The choices only mention DocumentDB.

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

There is no on-demand capacity mode for DocumentDB, though there is on-demand vCPU based pricing available.

upvoted 1 times

ninomfr64 1 year, 5 months ago

There is no on-demand capacity for DocumentDB, however Elastic Cluster option is provided "Elastic Clusters enables you to elastically scale your document database to handle millions of writes and reads, with petabytes of storage capacity" see https://aws.amazon.com/documentdb/faqs/#:~:text=to%20learn%20more.-,Elastic%20Clusters,-What%20is%20Amazon

upvoted 1 times

2aa2222 10 months, 3 weeks ago

DocumentDB does support on-demand capacity: https://aws.amazon.com/documentdb/pricing/#:~:text=On-demand%20instances%20let%20you%20pay%20per%20second%2C,and%20having%20to%20guess%20the%20correct%20capacity

upvoted 1 times

ProMax 1 year, 10 months ago

Selected Answer: C

Amazon DocumentDB does NOT have on-demand capacity mode, so its option C.

upvoted 3 times

ninomfr64 1 year, 5 months ago

There is no on-demand capacity for DocumentDB, however Elastic Cluster option is provided "Elastic Clusters enables you to elastically scale your document database to handle millions of writes and reads, with petabytes of storage capacity" see https://aws.amazon.com/documentdb/faqs/#:~:text=to%20learn%20more.-,Elastic%20Clusters,-What%20is%20Amazon

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: D

I was leaning towards Option C but "Appropriately sized instances" is vague since the question does not state the size of Mongo DB. On-demand instances serve the purpose here, they are offered by DocumentDB, see the link
https://aws.amazon.com/documentdb/pricing/

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

its a c

upvoted 2 times

easytoo 2 years ago

c-c-c-c-c-c-c-c

On-demand capacity mode as suggested in D may not provide the same level of high availability as multi-Availability Zone deployments. So it's c-c-c-c-c-c-c-c for me.

upvoted 2 times

SkyZeroZx 2 years ago

Selected Answer: C

See best practices for amazon documentdb - instance sizing in docs.
Addicionally there is no on-demand capacity mode.

upvoted 2 times

F_Eldin 2 years, 1 month ago

Selected Answer: C

DocumentDB does indeed support on-demand capacity mode (Contrary to what other users say here)
https://aws.amazon.com/blogs/database/running-spiky-workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/

but this mode is good for spikey workloads and does not address the high availablity requirement

upvoted 3 times

F_Eldin 2 years, 1 month ago

The correct link https://www.applytosupply.digitalmarketplace.service.gov.uk/g-cloud/services/743016963590682

upvoted 2 times

[Removed] 1 year, 7 months ago

The content mentioned in your link and the original comment are both mentioning things related to DynamoDB. Your link is even worse which is describing DynamoDB but say it is for DocumentDB. Please study hard

upvoted 1 times

leehjworking 2 years, 1 month ago

Selected Answer: C

See best practices for amazon documentdb - instance sizing in docs.

upvoted 1 times

Sarutobi 2 years, 2 months ago

Selected Answer: C

Going wit C. I still call the DocumentDB used in mode C "on-demand mode" because you have to select the Ec2 instance; the pricing documentation still uses that name. There is an Elastic cluster for DocumentDB. Could it be that option D "on-demand capacity mode" is referring to Elastic mode?

upvoted 2 times

OCHT 2 years, 2 months ago

Selected Answer: C

Amazon DocumentDB does not support an on-demand capacity mode. You can only choose from different instance classes that have fixed compute and memory resources. However, you can scale your instances up or down as needed, and you can also pause and resume your instances to save costs. Amazon DocumentDB also automatically scales your storage and I/O based on your data size and workload.

upvoted 1 times

Question #137 *Topic 1*

A digital marketing company has multiple AWS accounts that belong to various teams. The creative team uses an Amazon S3 bucket in its AWS account to securely store images and media files that are used as content for the company's marketing campaigns. The creative team wants to share the S3 bucket with the strategy team so that the strategy team can view the objects.

A solutions architect has created an IAM role that is named strategy_reviewer in the Strategy account. The solutions architect also has set up a custom AWS Key Management Service (AWS KMS) key in the Creative account and has associated the key with the S3 bucket. However, when users from the Strategy account assume the IAM role and try to access objects in the S3 bucket, they receive an Access Denied error.

The solutions architect must ensure that users in the Strategy account can access the S3 bucket. The solution must provide these users with only the minimum permissions that they need.

Which combination of steps should the solutions architect take to meet these requirements? (Choose three.)

A. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to the account ID of the Strategy account.

B. Update the strategy_reviewer IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.

C. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role.

D. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to an anonymous user.

E. Update the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role.

F. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.

Suggested Answer: *BCF*

Community vote distribution

ACF (100%)

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: ACF

B wrong - full permissions ? when question asks for minimum permissions.
D wrong - anonymous user ? anonymous does not work
E wrong - encrypt permissions ? No Strategy account needs decrypt permissions
So, A,C,F
upvoted 13 times

God_Is_Love 1 year, 9 months ago

first the source bucket needs to give grant access thru bucket policy and KMS key policy (A,C options)
Secondly, Strategy IAM role needs to give access to read from S3 bucket and also KMS key (Option F)
upvoted 3 times

leehjworking Highly Voted 1 year, 7 months ago

Selected Answer: ACF

B full permission ? X
D anonymous? X
E encryption not needed for strategy team
upvoted 6 times

career360guru Most Recent 1 year ago

Selected Answer: ACF

A, C and F
upvoted 1 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: ACF

By rule of elimination
BDE are wrong. God_Is_Love is spot on

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: ACF

its ACF

upvoted 2 times

OCHT 1 year, 8 months ago

Selected Answer: ACF

Option B suggests updating the strategy_reviewer IAM role to grant full permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key. This option is not ideal because it grants more permissions than necessary. The requirement is to provide users with only the minimum permissions they need to view objects in the S3 bucket.

Option D suggests creating a bucket policy that includes read permissions for the S3 bucket and setting the principal of the bucket policy to an anonymous user. This option is not ideal because it would allow anyone to read objects in the S3 bucket, which could pose a security risk.

Option E suggests updating the custom KMS key policy in the Creative account to grant encrypt permissions to the strategy_reviewer IAM role. This option is not necessary because the requirement is for users in the Strategy account to be able to view objects in the S3 bucket, not to encrypt them.

upvoted 3 times

mfsec 1 year, 9 months ago

Selected Answer: ACF

ACF is the best choice

upvoted 2 times

taer 1 year, 9 months ago

Selected Answer: ACF

A. Create a bucket policy that includes read permissions for the S3 bucket. Set the principal of the bucket policy to the account ID of the Strategy account.

C. Update the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role.

F. Update the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key.

upvoted 2 times

zozza2023 1 year, 11 months ago

Selected Answer: ACF

A C AND F

upvoted 3 times

Untamables 1 year, 11 months ago

Selected Answer: ACF

https://repost.aws/knowledge-center/cross-account-access-denied-error-s3

upvoted 3 times

masetromain 1 year, 11 months ago

Selected Answer: ACF

A, C, and F are the correct options.

upvoted 4 times

masetromain 1 year, 11 months ago

A, C, and F are the correct options.

Option A creates a bucket policy that includes read permissions for the S3 bucket and sets the principal of the bucket policy to the account ID of the Strategy account. This ensures that users in the Strategy account have the necessary permissions to access the S3 bucket.

Option C updates the custom KMS key policy in the Creative account to grant decrypt permissions to the strategy_reviewer IAM role. This ensures that the users in the Strategy account have the necessary permissions to decrypt the objects stored in the S3 bucket.

Option F updates the strategy_reviewer IAM role to grant read permissions for the S3 bucket and to grant decrypt permissions for the custom KMS key. This ensures that the users in the Strategy account have the necessary permissions to read the objects in the S3 bucket and to decrypt them using the custom KMS key.

The other options are not correct because they either grant unnecessary permissions (B, D) or grant permissions in the wrong way (E).

upvoted 3 times

zhangyu20000 1 year, 11 months ago

ACF is correct

upvoted 2 times

Question #138 *Topic 1*

A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data. Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days.

The company has a high-speed AWS Direct Connect connection. Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day.

Which solution meets these requirements?

A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS. When AWS receives the Snowball Edge device and the data is loaded into Amazon S3, use S3 events to trigger an AWS Lambda function to process the data.

B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3. Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data.

C. Use AWS DataSync to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow. Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data.

D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Batch job that executes on Amazon EC2 instances running the Docker containers to process the data.

Suggested Answer: *C*

Community vote distribution

C (76%) D (24%)

dev112233xx Highly Voted 2 years, 2 months ago

Selected Answer: C

Almost voted D because of the Storage Gateway + SAN combination.. but seems like it's not correct since S3 events cannot trigger Batch jobs directly, you need a Lambda function! S3 events can be only Lambda,SNS or SQS..

upvoted 23 times

Kampton 2 years, 2 months ago

Agree - The Lambda function acts as a bridge between the S3 event and AWS Batch, allowing you to trigger AWS Batch jobs in response to S3 events.

upvoted 3 times

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: D

Guys its Tricky one between C and D and answer is D! (Modernization question)
Look at this two below blogs :
https://aws.amazon.com/blogs/storage/using-aws-storage-gateway-to-modernize-next-generation-sequencing-workflows/

Thanks to tinyflame who made me do my research on this :-)
Yes, SAN -> Storage Gateway Only
NAS -> Data Sync or Storage Gateway
https://aws.amazon.com/blogs/storage/from-on-premises-to-aws-hybrid-cloud-architecture-for-network-file-shares/

upvoted 9 times

AWSum1 8 months, 3 weeks ago

Nope, you need S3 events to trigger Lambda. S3 events cannot trigger batch

upvoted 1 times

helloworldabc 10 months ago

just C

upvoted 1 times

God_Is_Love 2 years, 3 months ago

On Premise NAS and file servers to S3. --> Use DataSync solution

On Premise SMB or NFS file share to S3 --> Use Storage/File Gateway solution

upvoted 4 times

titi_r 1 year, 3 months ago

@God_Is_Love, both articles you've provided are NOT mentioning "SAN" at all. You cannot copy data from SAN using storage GW, but you do it with DataSync ran from within a server, which is connected to that SAN. Research more on what SAN is and how does it work :)

upvoted 1 times

FZA24 Most Recent 8 months, 2 weeks ago

Selected Answer: C

DataSync + Direct Connect

S3 => Lambda => SF

Docker => ECR => Batch

upvoted 1 times

k10training02 10 months, 2 weeks ago

lambda solo dura 900 segundos me voy por la D

upvoted 1 times

helloworldabc 10 months ago

just C

upvoted 1 times

trungtd 1 year, 1 month ago

Selected Answer: C

Currently, S3 events can only push to three different types of destinations:

SNS topic, SQS Queue, AWS Lamba.

You cannot directly trigger a Batch job by S3 Event

upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

A = 200GB very now and then doesn't need Snowball Edge

B = Data Pipeline is ETL and not suitable in hybrid scenarios

C = correct (DataSync does the job, also the app is already container based and it works well with Batch that is suited for HPC kind of workload - genomic sequencing is a typical HPC workload)

D = even tough Storage Gateway does the job you cannot directly trigger a AWS Batch job from an S3 event, you need either a Lambda in the middle or enable EventBrdige notification and create a rule that triggers the AWS Batch Job

upvoted 3 times

cox1960 1 year, 5 months ago

... "The main requirement is that the data needs to be accessible over the network in a file format like NFS that DataSync supports."

upvoted 1 times

cox1960 1 year, 5 months ago

C - Amazon Q says "While it does not directly support SAN (storage area network), you can use AWS DataSync to transfer data from files stored on a SAN volume to AWS storage services like Amazon S3."

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C is better option. Though D is also possible but as the jobs are already container based C would be better.

Question is not clear whether containers used on-premise are docker based containers.

upvoted 2 times

mosalahs 1 year, 6 months ago

Selected Answer: C

Data Transfer --- > Data Sync

Data Integration --- > Storage GW

Data Orchestration --- > Data Pipeline

upvoted 3 times

Maygam 1 year, 6 months ago

Selected Answer: C

D doesn't seem to be correct as AWS Batch is not a destination for AWS S3 events.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/notification-how-to-event-types-and-destinations.html
upvoted 2 times

uC6rW1aB 1 year, 9 months ago

Selected Answer: C

Option C: Use AWS DataSync to transfer data to Amazon S3. DataSync is designed for fast, easy and secure data transfer. This option also uses S3 events to trigger an AWS Lambda function, which launches an AWS Step Functions workflow and runs a Docker container using AWS Batch. This option takes into account data transfer, processing and container management, and should be the most suitable solution.
Option D: Use AWS Storage Gateway's file gateway to transfer data to Amazon S3. Storage Gateway is suitable for hybrid cloud environments, but in this case, since the company already has a high-speed AWS Direct Connect connection, it will be more efficient to use DataSync.
upvoted 2 times

Ganshank 1 year, 10 months ago

C.
Of the given options C is probably the closest. Step Functions can be used to model the workflow. D does not specify this. DataSync can be used to transfer data [https://docs.aws.amazon.com/datasync/latest/userguide/s3-cross-account-transfer.html].
upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: D

I choose D. My rationale - 200GB data for 1 genome sequence, Lets say DirectConnect is 1Gbps line, DataSync cannot efficiently transfer the data to get the processing under 1 day.
Agree with God_Is_Love's hypothesis
upvoted 1 times

vn_thanhtung 1 year, 10 months ago

S3 event can't trigger direct AWS Batch job. => C
upvoted 1 times

ninomfr64 1 year, 5 months ago

Assuming DX is 1Gbps, it takes about 27 minutes to transfer 200GB. also, I don't see how Storage Gateway can speedup things. My point is that here both DataSynch and Storage Gateway can di the job, but you cannot trigger Batch job directly from S3 object event. Thus C
upvoted 1 times

RGR21 1 year, 10 months ago

Does the AWS DataSync support SAN?
upvoted 1 times

ggrodskiy 1 year, 11 months ago

Correct D.
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C
D would be an option if using volume gateway and lambda to trigger batch
datasync dont need to support NAS. agent can copy off of NFS or SMB mount of the NAS drive.
upvoted 1 times

Question #139 *Topic 1*

A company runs a content management application on a single Windows Amazon EC2 instance in a development environment. The application reads and writes static content to a 2 TB Amazon Elastic Block Store (Amazon EBS) volume that is attached to the instance as the root device. The company plans to deploy this application in production as a highly available and fault-tolerant solution that runs on at least three EC2 instances across multiple Availability Zones.

A solutions architect must design a solution that joins all the instances that run the application to an Active Directory domain. The solution also must implement Windows ACLs to control access to file contents. The application always must maintain exactly the same content on all running instances at any given point in time.

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

A. Create an Amazon Elastic File System (Amazon EFS) file share. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to install the application, join the instance to the AD domain, and mount the EFS file share.

B. Create a new AMI from the current EC2 Instance that is running. Create an Amazon FSx for Lustre file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to join the instance to the AD domain and mount the FSx for Lustre file system.

C. Create an Amazon FSx for Windows File Server file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three instances. Implement a user data script to install the application and mount the FSx for Windows File Server file system. Perform a seamless domain join to join the instance to the AD domain.

D. Create a new AMI from the current EC2 instance that is running. Create an Amazon Elastic File System (Amazon EFS) file system. Create an Auto Scaling group that extends across three Availability Zones and maintains a minimum size of three Instances. Perform a seamless domain join to join the instance to the AD domain.

Suggested Answer: *B*

Community vote distribution

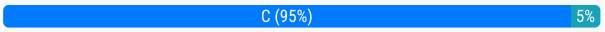

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: C

EFS is Linux/Mac based, So, A,D are out.
Lustre stands for Linux cluster, So B is out. Left is C which is correct (Amazon FSx for Windows)

upvoted 16 times

julmarcas Most Recent 8 months, 1 week ago

Selected Answer: C

C for windows, AD and ACLs

upvoted 1 times

rootcode 10 months, 3 weeks ago

Selected Answer: C

C is the correct option

upvoted 1 times

career360guru 1 year ago

Selected Answer: C

Option C as it is windows based OS.

upvoted 1 times

uC6rW1aB 1 year, 3 months ago

Selected Answer: C

Option B FSx for Lustre is not for Linux POSIX-compliant
Option C correct

upvoted 2 times

dkcloudguru 1 year, 3 months ago

C FSx for windows is a good fit for this
upvoted 1 times

Sam202 1 year, 5 months ago
FSx for Lustre can only be used by Linux-based instances.
upvoted 1 times

NikkyDicky 1 year, 5 months ago
Selected Answer: C
C for windows
upvoted 1 times

SkyZeroZx 1 year, 6 months ago
Selected Answer: C
EFS and FSx for Lustre == Linux
FSx Windows File == Windows
upvoted 3 times

mfsec 1 year, 9 months ago
Selected Answer: C
EFS and Windows is not straight forward. C is the best solution.
upvoted 2 times

zejou1 1 year, 9 months ago
Selected Answer: C
Amazon FSx is built on Windows Server... Access Control Lists (ACLs)... To control user access, Amazon FSx integrates with your on-premises Microsoft Active Directory as well as with AWS Microsoft Managed AD.
https://aws.amazon.com/fsx/windows/features/?nc=sn&loc=2

All others don't work - forget about the "least management" statement - it says "implement Windows ACLS to control..." all others are thrown out.
upvoted 3 times

kiran15789 1 year, 10 months ago
Selected Answer: C
Option D suggests using an EFS file system, which is a shared file system that can be mounted on multiple EC2 instances, but this requires additional configuration to keep the content in sync across all instances.

Option C is the optimal choice because Amazon FSx for Windows File Server supports Windows ACLs and seamlessly integrates with Active Directory to join instances to a domain. This option minimizes management overhead by reducing the complexity of managing multiple EFS file shares or writing scripts to synchronize content across EC2 instances.
upvoted 2 times

Musk 1 year, 10 months ago
Selected Answer: C
FSX for WIndows is the only option. The rest of options are not supported.
upvoted 2 times

jojom19980 1 year, 11 months ago
Selected Answer: C
FSx for Lustre can only be used by Linux-based instances.
upvoted 2 times

zozza2023 1 year, 11 months ago
Selected Answer: D
good answer are C or D but as it says LEAST management overhead ==> D as in C we will need a user data script
upvoted 1 times

zozza2023 1 year, 11 months ago
sorry D is uncorrect as it use Elastic File System (Amazon EFS) itch is not windows so Iswitch to C
upvoted 1 times

lxrdm 1 year, 5 months ago
Also that means each instance launched from the AMI will have 2TB EBS volume.. which is not ideal
upvoted 1 times

ARLV 1 year, 11 months ago

@masetromain is this a good exam study guide? Like how many questions were from here. Any help would be appreciated. Thank you

upvoted 1 times

Untamables 1 year, 11 months ago

Selected Answer: C

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

https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_join_instance.html

upvoted 1 times

Question #140 *Topic 1*

A software as a service (SaaS) based company provides a case management solution to customers A3 part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.

The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.

Which solution will meet these requirements MOST cost-effectively?

A. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.

B. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.

C. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in Amazon Simple Email Service (Amazon SES) with parameters for the customer data. Create an AWS Lambda function to call the SES template and to pass customer data to replace the parameters. Use the AWS Marketplace SMTP server to send the email message.

D. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template on Amazon SES with parameters for the customer data. Create an AWS Lambda function to call the SendTemplatedEmail API operation and to pass customer data to replace the parameters and the email destination.

Suggested Answer: *B*

Community vote distribution

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: D

SendTemplatedEmail
SendEmail
SendRawEmail are email api methods used in SES

upvoted 12 times

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: D

The correct answer is D.

In this solution, the company can use Amazon SES to send email messages, which will minimize operational overhead as SES is a fully managed service that handles sending and receiving email messages. The company can store the email template on Amazon SES with parameters for the customer data and use an AWS Lambda function to call the SendTemplatedEmail API operation, passing in the customer data to replace the parameters and the email destination. This solution eliminates the need to set up and manage an SMTP server on EC2 instances, which can be costly and time-consuming.

Option A and B are not correct because it requires to set up an SMTP server on EC2 instances, which is not necessary and will increase operational overhead.
Option C is not correct because it stores the email template in Amazon SES with parameters for the customer data which is not possible.

upvoted 11 times

Maria2023 1 year, 6 months ago

Ok, so according to chatgpt C is not correct because "Option C is not correct because it stores the email template in Amazon SES with parameters for the customer data which is not possible."
However, D says exactly the same - so D is not correct as well?
Do not fully trust chatgp

upvoted 7 times

titi_r 9 months, 2 weeks ago

ChatGPT also is saying "Option A and B are not correct because it requires to set up an SMTP server on EC2 instances", but those options are "A" and "C", not "A" and "B". Seems there is some mismatch with the options.
upvoted 2 times

[Removed] Most Recent 7 months, 3 weeks ago

Selected Answer: D

S3 Buckets is not needed to store template
upvoted 1 times

career360guru 1 year ago

Selected Answer: D

Option D
upvoted 1 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: D

D - Can send templated email with request parameters
upvoted 1 times

Jonalb 1 year, 5 months ago

Selected Answer: D

DDDDDDDD
upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: D

its a d
upvoted 1 times

Maria2023 1 year, 6 months ago

Selected Answer: B

I vote for B due to the fact that I cannot see an option to "Store the email template on Amazon SES with parameters for the customer data" Other than that it looks like a good option but it's just not working
upvoted 1 times

carpa_jo 1 year ago

D is correct.
Regarding your concerns about email templates on SES with parameters see: https://docs.aws.amazon.com/ses/latest/dg/send-personalized-email-api.html
upvoted 1 times

SK_Tyagi 1 year, 4 months ago

https://docs.aws.amazon.com/ses/latest/APIReference-V2/API_CreateEmailTemplate.html
upvoted 1 times

pk0619 6 months, 1 week ago

There can be variables in the template.
upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: D

keyword = SendTemplatedEmail API
upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: D

Template - easy one.
upvoted 1 times

zozza2023 1 year, 11 months ago

Selected Answer: D

D should be the answer
upvoted 3 times

zhangyu20000 1 year, 11 months ago

D is correct - https://docs.aws.amazon.com/ses/latest/APIReference/API_SendTemplatedEmail.html

upvoted 2 times

Question #141 *Topic 1*

A company is processing videos in the AWS Cloud by Using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.

The company has configured the SQS queue with a redrive policy that specifies a target dead-letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.

Several times during the day. the development team receives notification that messages are in the dead-letter queue and that videos have not been processed property. An investigation finds no errors m the application logs.

How can the company solve this problem?

A. Turn on termination protection tor the EC2 Instances

B. Update the visibility timeout for the SQS queue to 3 hours

C. Configure scale-in protection for the instances during processing

D. Update the redrive policy and set maxReceiveCount to 0.

Suggested Answer: *D*

Community vote distribution

C (70%) D (27%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: C

The correct answer is C. The company can solve the problem by configuring scale-in protection for the instances during processing. This will ensure that the instances are not terminated while they are processing videos. This will prevent the messages from moving to the dead-letter queue and ensure that videos are processed properly.

Option A is incorrect because turning on termination protection for the EC2 instances will not solve the problem as it will impact the ability of the Auto Scaling group to scale instances in and out based on the number of videos in the queue.

Option B is incorrect because the company has specified a visibility timeout of 1 hour, which is enough time for the instances to process a video and there is no need to update the timeout to 3 hours.

Option D is incorrect because the company has set the maxReceiveCount to 1 and changing it to 0 will not solve the problem. maxReceiveCount allowed range is 1 to 1000.

upvoted 28 times

[Removed] 1 year, 11 months ago

fully agree, option d is inocrrect because 0 is an invalida value for maxReceiveCount

upvoted 1 times

Bwitch 2 years, 1 month ago

ChatGPT confirms this reasoning.

upvoted 8 times

venvig Highly Voted 1 year, 10 months ago

Selected Answer: C

Refer https://aws.amazon.com/blogs/aws/new-instance-protection-for-auto-scaling/

From the above link, "an instance might be handling a long-running work task, perhaps pulled from an SQS queue. Protecting the instance from termination will avoid wasted work" - This is what the question is also alluding to.

This is how one would make use of the functionality.

You change the protection status of one or more instances by calling the SetInstanceProtection function. If you wanted to use this function to protect long-running, queue-driven worker processes from scale-in termination, you could set up your application as follows (this is pseudocode):

while (true)

{
SetInstanceProtection(False);
Work = GetNextWorkUnit();
SetInstanceProtection(True);
ProcessWorkUnit(Work);
SetInstanceProtection(False);
}
upvoted 6 times

Jorkaef Most Recent 7 months, 2 weeks ago
Correct is C:

B. 3-hour visibility timeout

Too long for 30-minute processing
Could delay reprocessing of failed messages
Doesn't address root cause

C. Scale-in protection during processing

Prevents instance termination while processing
Allows message processing to complete
Prevents message return to queue
Stops premature scale-in
✓ CORRECT
upvoted 1 times

Jorkaef 7 months, 2 weeks ago
B is correct;
updating the visibility timeout to 3 hours (option B) is the most appropriate solution as it gives enough time for the messages to be processed without being prematurely marked as failures.
upvoted 1 times

trungtd 1 year ago
Selected Answer: D
D is a typo
upvoted 1 times

VerRi 1 year, 4 months ago
Selected Answer: D
If Option D is a typo, then D
upvoted 2 times

Greanny 1 year, 5 months ago
B.
The best solution for this problem is to update the visibility timeout for the SQS queue to 3 hours. This is because when the visibility timeout is set to 1 hour, it means that if the EC2 instance doesn't process the message within an hour, it will be moved to the dead-letter queue. By increasing the visibility timeout to 3 hours, this should give the EC2 instance enough time to process the message before it gets moved to the dead-letter queue. Additionally, configuring scale-in protection for the EC2 instances during processing will help to ensure that the instances are not terminated while the messages are being processed.
upvoted 3 times

tmlong18 1 year, 5 months ago
Selected Answer: D
Option D is a typo.
I seen the same question in udemy but the Option D is 10
upvoted 4 times

career360guru 1 year, 6 months ago
Selected Answer: C
Option C is correct.
upvoted 2 times

severlight 1 year, 7 months ago

Selected Answer: C

setting MaxReceiveCount to 0 doesn't make and send and it impossible, because messages would be send to DLQ without any attempt to consume them from source queue

upvoted 1 times

Russs99 1 year, 9 months ago

Selected Answer: D

checked 4 AI, C is definitely not the correct answer: Option C: Configuring scale-in protection for the instances during processing will not prevent messages from being moved to the dead-letter queue if they cannot be processed on the first attempt.

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: C

Going with C only because D has value of maxReceiveCount set to 0

upvoted 2 times

rtguru 1 year, 11 months ago

I go with C

upvoted 1 times

YodaMaster 1 year, 12 months ago

Selected Answer: B

B.

AWS "recommends setting your queue's visibility timeout to six times your function timeout" which makes 3 hours perfect.

source: https://docs.aws.amazon.com/lambda/latest/dg/with-sqs.html

upvoted 2 times

ajeeshb 1 year, 3 months ago

But this for a queue to use with lambda. Here it is EC2 in ASG

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C more likely

upvoted 1 times

Maria2023 2 years ago

I couldn't find any way to configure scale-in protection for the instances during processing except to do it manually, which is going to be an insane exercise. Eventually, that can be done by the application as part of the processing but I would then expect some more context in the answer.

upvoted 1 times

dev112233xx 2 years, 1 month ago

Selected Answer: D

D makes sense

I think D answer has a typo! probably they didn't copy the text properly

https://repost.aws/knowledge-center/lambda-retrying-valid-sqs-messages

upvoted 6 times

Question #142 *Topic 1*

A company has developed APIs that use Amazon API Gateway with Regional endpoints. The APIs call AWS Lambda functions that use API Gateway authentication mechanisms. After a design review, a solutions architect identifies a set of APIs that do not require public access.

The solutions architect must design a solution to make the set of APIs accessible only from a VPC. All APIs need to be called with an authenticated user

Which solution will meet these requirements with the LEAST amount of effort?

A. Create an internal Application Load Balancer (ALB). Create a target group. Select the Lambda function to call. Use the ALB DNS name to call the API from the VPC.

B. Remove the DNS entry that is associated with the API in API Gateway. Create a hosted zone in Amazon Route 53. Create a CNAME record in the hosted zone. Update the API in API Gateway with the CNAME record. Use the CNAME record to call the API from the VPC.

C. Update the API endpoint from Regional to private in API Gateway. Create an interface VPC endpoint in the VPCreate a resource policy, and attach it to the API. Use the VPC endpoint to call the API from the VPC.

D. Deploy the Lambda functions inside the VPC Provision an EC2 instance, and install an Apache server. From the Apache server, call the Lambda functions. Use the internal CNAME record of the EC2 instance to call the API from the VPC.

Suggested Answer: *D*

Community vote distribution

C (100%)

bjexamprep Highly Voted 1 year, 6 months ago

Selected Answer: C

Bad question design. None of the answers is correct.
None of the answers mentions how to satisfy the requirement of "All APIs need to be called with an authenticated user".
Another requirement "make the set of APIs accessible only from a VPC". "the set" doesn't mean the whole set. Here "the set" means a part of the whole set.
A: The set of APIs are still publicly accessible.
B: Removing DNS entry doesn't remove the public accessibility.
C: This is making the whole set of APIs private. If this answer can be specific to "the set" APIs, this could be a good answer.
D: Using EC2 instances is always a bad answer.
upvoted 11 times

altonh 5 months, 2 weeks ago

Agree. The proper solution should be:
Create a new private API GW and move those private APIs to this newly created API GW.
upvoted 1 times

toma 1 year ago

there is only set of APIs that do not require public access, you dont need all APIs private access? so it could be that the answer is A?
upvoted 2 times

zozza2023 Highly Voted 2 years, 5 months ago

Selected Answer: C

should be C as on the question has said 'no need for public IP" ==> private in API gateway = VPC endpoint
upvoted 9 times

AimarLeo Most Recent 1 year, 5 months ago

All given answers are not ideal.. the closet one is C BUT.. .when mentioning the requirement to have only 'a set of API to be private' means 'not all'.. turning the endpoint from public to private will turn all to Private ,, which is not fully correct as per the question.. I suppose the given answer or question missing an info.. or AWS starts playing with AI
upvoted 3 times

carpa_jo 1 year, 6 months ago

Selected Answer: C

https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-private-apis.html

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C

upvoted 1 times

venvig 1 year, 10 months ago

Selected Answer: C

Refer https://aws.amazon.com/blogs/compute/introducing-amazon-api-gateway-private-endpoints/

upvoted 1 times

Explorer_30 1 year, 10 months ago

Answer is C as explain in https://repost.aws/knowledge-center/api-gateway-vpc-connections

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: C

Regional to Private fits the use-case

upvoted 1 times

rtguru 1 year, 11 months ago

the best possible answer from all the options is C

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

it's C, although it begs the questions about APIs that need to stay public...

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: C

C. Update the API endpoint from Regional to private in API Gateway.

upvoted 1 times

masetromain 2 years, 5 months ago

Selected Answer: C

The correct answer is C. Update the API endpoint from Regional to private in API Gateway. Create an interface VPC endpoint in the VPC. Create a resource policy, and attach it to the API. Use the VPC endpoint to call the API from the VPC.

This solution will meet the requirements with the least amount of effort because it utilizes the built-in features of API Gateway and VPC to restrict access to the API. With this method, no additional infrastructure or configurations are necessary.

A and B are not correct because they would require additional infrastructure and configurations.

D is not correct because it would require provisioning an EC2 instance and installing an Apache server, introducing additional complexity and management overhead.

upvoted 5 times

zhangyu20000 2 years, 5 months ago

C is correct

upvoted 1 times

Question #143 *Topic 1*

A weather service provides high-resolution weather maps from a web application hosted on AWS in the eu-west-1 Region. The weather maps are updated frequently and stored in Amazon S3 along with static HTML content. The web application is fronted by Amazon CloudFront.

The company recently expanded to serve users in the us-east-1 Region, and these new users report that viewing their respective weather maps is slow from time to time.

Which combination of steps will resolve the us-east-1 performance issues? (Choose two.)

A. Configure the AWS Global Accelerator endpoint for the S3 bucket in eu-west-1. Configure endpoint groups for TCP ports 80 and 443 in us-east-1.

B. Create a new S3 bucket in us-east-1. Configure S3 cross-Region replication to synchronize from the S3 bucket in eu-west-1.

C. Use Lambda@Edge to modify requests from North America to use the S3 Transfer Acceleration endpoint in us-east-1.

D. Use Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1.

E. Configure the AWS Global Accelerator endpoint for us-east-1 as an origin on the CloudFront distribution. Use Lambda@Edge to modify requests from North America to use the new origin.

Suggested Answer: *BD*

Community vote distribution

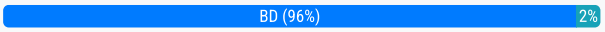

sambb Highly Voted 2 years, 3 months ago

Selected Answer: BD

A: Global Accelerator can't have an s3 bucket as endpoint
C: People are complaining about time to retreive maps. Transfert acceleration is used to accelerate PUT requests to an s3 bucket located in a distant region.
E: An accelerator as cloudfront origin does not make much sense, because cloudfront is already using the AWS network. Global Accelerator is usually for Layer 4 networking and/or static anycast IPs

upvoted 19 times

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: BD

B is correct because it involves creating a new S3 bucket in the us-east-1 region and configuring cross-Region replication to synchronize from the existing S3 bucket in eu-west-1. This will allow users in us-east-1 to access the weather maps from a closer location, improving performance.

D is correct because it involves using Lambda@Edge to modify requests from North America to use the S3 bucket in us-east-1. This will also allow users in us-east-1 to access the weather maps from a closer location, improving performance.

A and E are not correct because they do not involve creating a new S3 bucket in us-east-1, which is necessary for improving performance for the users in that region. C is not correct because it involves using the S3 Transfer Acceleration endpoint, which is a different service and not necessary for this scenario.

upvoted 8 times

ahhatem Most Recent 6 months, 2 weeks ago

Selected Answer: BD

Although, D is not really correct. You should be using "s3 multi-region access point". It is designed specifically for this scenario.

upvoted 1 times

altonh 5 months, 2 weeks ago

The correct answer implies a CloudFront with multiple origins, i.e. pointing to two (2) S3 buckets and using Lambda@Edge to decide which origin to go to.

upvoted 1 times

pangchn 1 year, 2 months ago

Selected Answer: BD

BD

C using S3 Transfer Acceleration is good but this answer option itself is wrong due to the statement that pointing to a regional endpoint, where it doesn't exist. Once enable, it is just a global endpoint URL

https://docs.aws.amazon.com/AmazonS3/latest/userguide/transfer-acceleration-examples.html

upvoted 1 times

jpa8300 1 year, 5 months ago

Selected Answer: AC

If you want to improve latency , you always look for Global Accelerator fro the readings and Transfer accelerator for the updates.

Yes, it is possible to configure AWS Global Accelerator to distribute traffic from an S3 bucket in one AWS Region (eu-west-1) to endpoint groups in another AWS Region (us-east-1) for TCP ports 80 and 443. This configuration can be useful for improving the performance and availability of your S3 bucket for users in both regions.

This way you sabe money in the storage, you don't need to duplicate the storage. And for persons that chose option D, if you update the bucket there, those objects will not be replicated to the other region since replication works only in one way.

upvoted 1 times

helloworldabc 10 months ago

just BD

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: BD

Option B & D

upvoted 1 times

bjexamprep 1 year, 6 months ago

Selected Answer: BD

This is not a good question design. Does that mean the application use CloudFront in EU and does not use CloudFront in the US? How weird it is!!!

upvoted 3 times

Jrhp 1 year, 7 months ago

Selected Answer: BD

Exactly case from this blog post https://aws.amazon.com/blogs/networking-and-content-delivery/dynamically-route-viewer-requests-to-any-origin-using-lambdaedge/

upvoted 4 times

rtguru 1 year, 11 months ago

BD, I was initially looking at BE, I think global accelerator is used more for write requests.

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: BD

BD makes more ense

upvoted 2 times

SmileyCloud 1 year, 12 months ago

Selected Answer: BD

https://godof.cloud/dynamic-origin-s3-spa/

Use case

upvoted 1 times

Eshu2009 2 years, 3 months ago

BE- global accelerators improve performance by providing edge location for onboarding traffic.

upvoted 3 times

Eshu2009 2 years, 3 months ago

Q: Can I use AWS Global Accelerator for object storage with Amazon S3?

A: You can use Amazon S3 Multi-Region Access Points to get the benefits of Global Accelerator for object storage. S3 Multi-Region Access Points use Global Accelerator transparently to provide a single global endpoint to access a data set that spans multiple S3 buckets in different AWS Regions. This allows you to build multi-region applications with the same simple architecture used in a single region, and then to run those applications anywhere in the world. Application requests made to an S3 Multi-Region Access Point's global endpoint automatically route over the AWS global network to the S3 bucket with the lowest network latency. This allows applications to automatically avoid congested network segments on the public internet, improving application performance and reliability.

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: BD

Ill go with BD

upvoted 1 times

kiran15789 2 years, 4 months ago

Selected Answer: BD

Since only one additional region we dont need global accelerators

upvoted 4 times

bititan 2 years, 4 months ago

Selected Answer: BC

S3 transfer acceleration is more efficient

upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: BD

A and E are not correct as there isn't a need to use aws global accel

upvoted 2 times

zhangyu20000 2 years, 5 months ago

BD is correct

upvoted 1 times

Question #144 *Topic 1*

A solutions architect is investigating an issue in which a company cannot establish new sessions in Amazon Workspaces. An initial analysis indicates that the issue involves user profiles. The Amazon Workspaces environment is configured to use Amazon FSx for Windows File Server as the profile share storage. The FSx for Windows File Server file system is configured with 10 TB of storage.

The solutions architect discovers that the file system has reached Its maximum capacity. The solutions architect must ensure that users can regain access. The solution also must prevent the problem from occurring again.

Which solution will meet these requirements?

A. Remove old user profiles to create space. Migrate the user profiles to an Amazon FSx for Lustre file system.

B. Increase capacity by using the update-file-system command. Implement an Amazon CloudWatch metric that monitors free space. Use Amazon EventBridge to invoke an AWS Lambda function to increase capacity as required.

C. Monitor the file system by using the FreeStorageCapacity metric in Amazon CloudWatch. Use AWS Step Functions to increase the capacity as required.

D. Remove old user profiles to create space. Create an additional FSx for Windows File Server file system. Update the user profile redirection for 50% of the users to use the new file system.

Suggested Answer: *C*

Community vote distribution

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: B

B is correct. It can prevent the issue from happening again by monitoring the file system with the FreeStorageCapacity metric in Amazon CloudWatch and using Amazon EventBridge to invoke an AWS Lambda function to increase the capacity as required. This ensures that the file system always has enough free space to store user profiles and avoids reaching maximum capacity.
A: Removing old user profiles may not be sufficient to create enough space and does not prevent the problem from happening again.
C: AWS Step Functions cannot be used to increase capacity, it is a service for creating and running workflows that stitch together multiple AWS services.
D: Creating an additional FSx for Windows File Server file system and updating user profile redirection for a portion of the users may not be sufficient to prevent the problem from happening again and does not address the current capacity issue.

upvoted 8 times

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: B

https://docs.aws.amazon.com/cli/latest/reference/fsx/update-file-system.html
EventBridge invoking lambda to update settings will prevent too from occurring again

upvoted 8 times

sse69 Most Recent 7 months, 1 week ago

Selected Answer: B

Wouldn't you need a cloudwatch alarm that would trigger a Lambda based on the metric going above a certain treshold?
Metric -> Lambda is a bit of a shortcut

upvoted 1 times

red_panda 8 months ago

Selected Answer: D

It's D.
Option B Simply do not prevent problem to happen again. It's not possible to resize the FSx Size after creation so option D is more suitable.

upvoted 1 times

career360guru 1 year ago

Selected Answer: B

Option B

upvoted 1 times

rtguru 1 year, 5 months ago

B is the correct answer

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

it's B

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: B

keyword == update-file-system

upvoted 1 times

leehjworking 1 year, 7 months ago

Selected Answer: C

Is it necessary to implement new cloudwatch metric? And using step functions seems to be able to increase storage capacity, according to the following reference.

https://docs.aws.amazon.com/step-functions/latest/dg/supported-services-awssdk.html#supported-services-awssdk-list

upvoted 1 times

Maria2023 1 year, 6 months ago

Perhaps the metric is used to trigger the step functions

upvoted 1 times

OCHT 1 year, 8 months ago

Selected Answer: D

B. Increasing capacity using the update-file-system command is not applicable to FSx for Windows File Server. The command is for Amazon EFS, not FSx for Windows File Server.

upvoted 2 times

rbm2023 1 year, 7 months ago

StorageCapacity

Use this parameter to increase the storage capacity of an FSx for Windows File Server, FSx for Lustre, FSx for OpenZFS, or FSx for ONTAP file system. Specifies the storage capacity target value, in GiB, to increase the storage capacity for the file system that you're updating.

https://docs.aws.amazon.com/fsx/latest/APIReference/API_UpdateFileSystem.html

Example using the CLI

aws fsx update-file-system --file-system-id fs-0123456789abcdef0 --storage-capacity 10240

upvoted 5 times

yama234 1 year, 8 months ago

B

As you need additional storage, you can increase the storage capacity that is configured on your FSx for Windows File Server file system. You can do so using the Amazon FSx console, the Amazon FSx API, or the AWS Command Line Interface (AWS CLI).

upvoted 3 times

Cloud_noob 1 year, 8 months ago

Selected Answer: B

https://chat.openai.com/chat

upvoted 2 times

mfsec 1 year, 9 months ago

Selected Answer: B

B is correct

upvoted 2 times

zozza2023 1 year, 11 months ago

Selected Answer: B

B seems to be the correct answer.

the unique possible solution is to add storage capacity using CLI

upvoted 4 times

pitakk 1 year, 11 months ago

Selected Answer: B

To increase the storage capacity for an FSx for Windows File Server file system, use the AWS CLI command update-file-system.
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/managing-storage-capacity.html It's B.

upvoted 3 times

zhangyu20000 1 year, 11 months ago

B is correct. It can prevent issue happen again with EventBridge and Lambda
A: not make sense at all
C: Cannot use Step Function to increase capacity
D: not prevent happen again

upvoted 2 times

Question #145 *Topic 1*

An international delivery company hosts a delivery management system on AWS. Drivers use the system to upload confirmation of delivery. Confirmation includes the recipient's signature or a photo of the package with the recipient. The driver's handheld device uploads signatures and photos through FTP to a single Amazon EC2 instance. Each handheld device saves a file in a directory based on the signed-in user, and the file name matches the delivery number. The EC2 instance then adds metadata to the file after querying a central database to pull delivery information. The file is then placed in Amazon S3 for archiving.

As the company expands, drivers report that the system is rejecting connections. The FTP server is having problems because of dropped connections and memory issues in response to these problems, a system engineer schedules a cron task to reboot the EC2 instance every 30 minutes. The billing team reports that files are not always in the archive and that the central system is not always updated.

A solutions architect needs to design a solution that maximizes scalability to ensure that the archive always receives the files and that systems are always updated. The handheld devices cannot be modified, so the company cannot deploy a new application.

Which solution will meet these requirements?

A. Create an AMI of the existing EC2 instance. Create an Auto Scaling group of EC2 instances behind an Application Load Balancer. Configure the Auto Scaling group to have a minimum of three instances.

B. Use AWS Transfer Family to create an FTP server that places the files in Amazon Elastic File System (Amazon EFS). Mount the EFS volume to the existing EC2 instance. Point the EC2 instance to the new path for file processing.

C. Use AWS Transfer Family to create an FTP server that places the files in Amazon S3. Use an S3 event notification through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function. Configure the Lambda function to add the metadata and update the delivery system.

D. Update the handheld devices to place the files directly in Amazon S3. Use an S3 event notification through Amazon Simple Queue Service (Amazon SQS) to invoke an AWS Lambda function. Configure the Lambda function to add the metadata and update the delivery system.

Suggested Answer: *B*

Community vote distribution

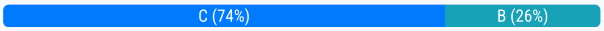

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: C

C is correct. Using AWS Transfer Family to create an FTP server that places the files in Amazon S3 and using S3 event notifications through Amazon Simple Notification Service (Amazon SNS) to invoke an AWS Lambda function will ensure that the archive always receives the files and that the central system is always updated. This solution maximizes scalability and eliminates the need for manual intervention, such as rebooting the EC2 instance.

Option A and B still use EC2 instance, which is the source of the problem. Option D requires modification to the handheld devices which is not possible.

upvoted 14 times

venvig Highly Voted 1 year, 3 months ago

Selected Answer: B

I agree that "C" is the ideal design.
But here the question states that :
Ec2 instance is running the SFTP server.
File is uploaded from handheld devices to a file system in the Ec2 instance.
The Ec2 instance then adds metadata to the file.
The file is then placed in s3.
The condition states that:
The company cannot deploy a new application.

Based on the condition, if I use lambda to add meta data, then its like deploying a new application.
(We don't know if the application can be seamlessly rewritten in lambda. Will it finish under 15 mins ? etc.,)

If we strictly interpret this as not being able to introduce any new logic or components (like a Lambda function for metadata processing), then Option (B) is the answer.
Option B essentially replaces the FTP server with AWS Transfer Family and uses Amazon EFS as the file storage, which can scale and handle more connections. The existing EC2 instance, which already has the logic for metadata addition, would simply point to this new file path on EFS. This minimizes changes to the existing application logic.
upvoted 7 times

pk0619 6 months, 1 week ago
they had to reboot the ec2 because of memory, without scaling EC2 they will still have that problem and since B does nothing about adding more memory, it cannot be right choice.
upvoted 1 times

pk0619 6 months, 1 week ago
Actually offloading FTP from EC2 might eliminate memory issue, so it could very well be B as well
upvoted 1 times

gofavad926 9 months, 2 weeks ago
the text is: "The handheld devices cannot be modified, so the company cannot deploy a new application". Following your comment, you can't use neither the AWS Transfer Family. This is also new :D
upvoted 2 times

kgcain 1 year, 2 months ago
From the app description, I am sure that it should work under 15min.
upvoted 1 times

EApeer Most Recent 9 months, 1 week ago
B is the best answer. The system is such that each handheld device saves a file in a directory based on the signed-in user, and the file name matches the delivery number. This means that we need a file storage that the data are stored hierarchically in a top-down network of folders. And a file system that has adaptive throughtput to resolve the dropped connections and memory issues. EFS will be the suitable solution component. S3 however has all the data stored on the same flat plane requiring more comprehensive metadata (labels) to make it manageable.
upvoted 1 times

kz407 9 months, 1 week ago
Selected Answer: C
It says "so the company cannot deploy a new application".

This means that it's the handheld devices they can't deploy a new application into. While B works, It still relies on one EC2 instance, which is a part of the problem.
upvoted 1 times

gofavad926 9 months, 2 weeks ago
Selected Answer: C
C, transfer family + S3
upvoted 1 times

zanhsieh 11 months ago
Selected Answer: C
C.
A: No. FTP is not HTTP / HTTPS. FTP -> NLB. HTTP / HTTPS -> ALB.
B: No. This needs extra steps (DataSync?) to move to S3, and the billing team would still complain about not always updated since it will be certain lag-behind time.
C: Correct.
D: No. S3 event notification can directly trigger Lambda.
upvoted 2 times

JMAN1 12 months ago
Selected Answer: C
C. does not require handheld device to be changed. And it solves EC2 dropped Conection by uisng S3.
upvoted 1 times

career360guru 1 year ago
Selected Answer: C
Option C
upvoted 1 times

Chung234 1 year, 2 months ago

Selected Answer: B

The answer is A.

Q: Can I use FTP with an internet-facing endpoint?

A: No, when you enable FTP, you will only be able to use VPC hosted endpoint's internal access option. If traffic needs to traverse the public network, secure protocols such as SFTP or FTPS should be used.

Source: https://aws.amazon.com/aws-transfer-family/faqs/

upvoted 3 times

ele 10 months ago

ALB is a load balancer that operates at Layer 7. Only HTTP and HTTPS can be used as ALB protocols.
Therefore, it is not possible to set ALB at the front of the FTP server.

upvoted 1 times

rtguru 1 year, 5 months ago

This one of those tricky questions. I'm not sure if to go with A or C

upvoted 1 times

rrrrrrrrr1 1 year, 5 months ago

IDK yall, it does say clearly "cannot deploy a new application" and the only instance of that is A.

I Agree C is better but IDK the semantics here

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

its a c

upvoted 1 times

Maria2023 1 year, 6 months ago

Selected Answer: C

Since AWS Transfer Family supports Amazon S3 Access Point then it's a standard scenario - FTP->S3->Event->Lambda. Scalable and serverless

upvoted 2 times

Jackhemo 1 year, 6 months ago

Selected Answer: C

olabiba.ai says C.

1. Scalability: By using AWS Transfer Family to create an FTP server that places the files directly in Amazon S3, you can leverage the scalability and durability of S3. S3 is designed to handle high volumes of data and can scale seamlessly as your company expands.

2. Reliability: With S3 as the destination for the files, you can ensure that the archive always receives the files. S3 provides high durability and availability, reducing the chances of data loss.

3. System updates: By using an S3 event notification through Amazon SNS, you can trigger an AWS Lambda function whenever a new file is uploaded to S3. This Lambda function can then add the necessary metadata and update the delivery system, ensuring that the central system is always updated.

4. No modification to handheld devices: Since the handheld devices cannot be modified, this solution allows the devices to continue uploading files through FTP. The only change is the destination, which is now the S3 bucket.

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: C

C is the most efficient

upvoted 3 times

zozza2023 1 year, 11 months ago

Selected Answer: C

C is correct

upvoted 3 times

zhangyu20000 1 year, 11 months ago

C is correct

upvoted 2 times

Question #146 *Topic 1*

A company is running an application in the AWS Cloud. The application runs on containers m an Amazon Elastic Container Service (Amazon ECS) cluster. The ECS tasks use the Fargate launch type. The application's data is relational and is stored in Amazon Aurora MySQL. To meet regulatory requirements, the application must be able to recover to a separate AWS Region in the event of an application failure. In case of a failure, no data can be lost.

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

A. Provision an Aurora Replica in a different Region.

B. Set up AWS DataSync for continuous replication of the data to a different Region.

C. Set up AWS Database Migration Service (AWS DMS) to perform a continuous replication of the data to a different Region.

D. Use Amazon Data Lifecycle Manager (Amazon DLM) to schedule a snapshot every 5 minutes.

Suggested Answer: *A*

Community vote distribution

A (100%)

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: A

A is correct. Provision an Aurora Replica in a different Region will meet the requirement of the application being able to recover to a separate AWS Region in the event of an application failure, and no data can be lost, with the least amount of operational overhead.

B. AWS DataSync can replicate data, but it is not a fully managed service and requires more configuration and management.

C. AWS DMS is a fully managed service for migrating data between databases, but it may require additional configuration and management to continuously replicate data in real-time.

D. Amazon DLM can be used for scheduling snapshots, but it does not provide real-time replication and may not meet the requirement of no data loss in case of a failure.

upvoted 8 times

career360guru Most Recent 1 year ago

Selected Answer: A

Option A

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

its an A

upvoted 2 times

Goatin 1 year, 7 months ago

When you provision an Aurora Replica in a different AWS Region, the replica is kept in sync with the primary database using Aurora's replication capabilities. In the event of a failure in the primary Region, you can promote the Aurora Replica to become the new primary database, which allows you to continue operations with no data loss.

However, provisioning and maintaining an Aurora Replica in a different AWS Region requires ongoing management and monitoring to ensure that it stays in sync with the primary database

upvoted 3 times

mfsec 1 year, 9 months ago

Selected Answer: A

Replica

upvoted 4 times

God_Is_Love 1 year, 9 months ago

Selected Answer: A

B,C are on premises usecase solutions. D is wrong because 5 minute worth of data could be lost against the requirement. So A is correct. In fact replica works as standby if primary DB fails.

upvoted 4 times

zozza2023 1 year, 11 months ago

Selected Answer: A

A is correct

upvoted 4 times

zhangyu20000 1 year, 11 months ago

A is correct

B: cannot use DataSync for Aurora backup

C: too complex

D: DLM is for EBS backup. Here use managed Aurora server, no access to EBS

upvoted 2 times

Question #147 *Topic 1*

A financial services company receives a regular data feed from its credit card servicing partner. Approximately 5,000 records are sent every 15 minutes in plaintext, delivered over HTTPS directly into an Amazon S3 bucket with server-side encryption. This feed contains sensitive credit card primary account number (PAN) data. The company needs to automatically mask the PAN before sending the data to another S3 bucket for additional internal processing. The company also needs to remove and merge specific fields, and then transform the record into JSON format. Additionally, extra feeds are likely to be added in the future, so any design needs to be easily expandable.

Which solutions will meet these requirements?

A. Invoke an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Invoke another Lambda function when new messages arrive in the SQS queue to process the records, writing the results to a temporary location in Amazon S3. Invoke a final Lambda function once the SQS queue is empty to transform the records into JSON format and send the results to another S3 bucket for internal processing.

B. Invoke an AWS Lambda function on file delivery that extracts each record and writes it to an Amazon SQS queue. Configure an AWS Fargate container application to automatically scale to a single instance when the SQS queue contains messages. Have the application process each record, and transform the record into JSON format. When the queue is empty, send the results to another S3 bucket for internal processing and scale down the AWS Fargate instance.

C. Create an AWS Glue crawler and custom classifier based on the data feed formats and build a table definition to match. Invoke an AWS Lambda function on file delivery to start an AWS Glue ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, have the ETL job send the results to another S3 bucket for internal processing.

D. Create an AWS Glue crawler and custom classifier based upon the data feed formats and build a table definition to match. Perform an Amazon Athena query on file delivery to start an Amazon EMR ETL job to transform the entire record according to the processing and transformation requirements. Define the output format as JSON. Once complete, send the results to another S3 bucket for internal processing and scale down the EMR cluster.

Suggested Answer: *C*

Community vote distribution

C (100%)

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: C

Extract Data from S3 + mask + Send to another S3 + Transform/Process + Load into S3
All these are ETL, ELT tasks which should ring Glue

EMR is more focused on big data processing frameworks such as Hadoop and Spark,
while Glue is more focused on ETL, More over 5000 records every 15 minutes is not soo big data..So I choose C
upvoted 21 times

tycho 1 year, 8 months ago

EMR and Glue are the same; Glue is managed cluster by AWS , EMR customer manages the clutster
upvoted 2 times

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: C

C is correct. It will process the data in batch mode using Glue ETL job which can handle large amount of data and can be scheduled to run periodically. This solution is also easily expandable for future feeds.

A: It uses multiple Lambda functions, SQS queue and S3 temporary location which will increase operational overhead.
B: Using Fargate may not be the most cost-effective solution and also it may not handle large amount of data.
D: Athena and EMR both are powerful tools but they are more complex and can be more costly than Glue.
upvoted 7 times

career360guru Most Recent 1 year ago

Selected Answer: C

Option C

upvoted 1 times

totten 1 year, 2 months ago

Selected Answer: C

Option C is the most suitable solution for the described scenario:

1) AWS Glue Crawler and Custom Classifier: Use AWS Glue to create a crawler and custom classifier to understand and catalogue the data feed formats. This step ensures that AWS Glue can work with the incoming data effectively.

2) AWS Glue ETL Job: Create an AWS Lambda function that triggers an AWS Glue ETL job when a new data file is delivered. This ETL job can perform the required transformation, including masking, field removal, and converting records to JSON format. AWS Glue is a suitable service for data preparation and transformation.

3) Output to S3 Bucket.

This approach is scalable, easily expandable to handle additional feeds in the future, and leverages AWS Glue's capabilities for data transformation and processing. It also maintains a clear separation of tasks, making it a robust and efficient solution for the given requirements.

upvoted 3 times

dkcloudguru 1 year, 3 months ago

C is the good option EMR(Big data, Spark, Hadoop) is for near real-time data processing and it isn't a good fit in this case

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

its a C

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: C

EMR is big data but not is need in this case

then AWS Glue + Lambdas + S3 is good option

C

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: C

C makes the most sense.

upvoted 2 times

Musk 1 year, 10 months ago

The question is at what point Athena and EMR are a better choice because it is a lot of data to store and process

upvoted 1 times

Sarutobi 1 year, 9 months ago

That, I agree. Honestly, I will use it from day one, regardless.

upvoted 1 times

zozza2023 1 year, 11 months ago

Selected Answer: C

C is correct.

upvoted 4 times

zhangyu20000 1 year, 11 months ago

C is correct

upvoted 1 times

Question #148 *Topic 1*

A company wants to use AWS to create a business continuity solution in case the company's main on-premises application fails. The application runs on physical servers that also run other applications. The on-premises application that the company is planning to migrate uses a MySQL database as a data store. All the company's on-premises applications use operating systems that are compatible with Amazon EC2.

Which solution will achieve the company's goal with the LEAST operational overhead?

A. Install the AWS Replication Agent on the source servers, including the MySQL servers. Set up replication for all servers. Launch test instances for regular drills. Cut over to the test instances to fail over the workload in the case of a failure event.

B. Install the AWS Replication Agent on the source servers, including the MySQL servers. Initialize AWS Elastic Disaster Recovery in the target AWS Region. Define the launch settings. Frequently perform failover and fallback from the most recent point in time.

C. Create AWS Database Migration Service (AWS DMS) replication servers and a target Amazon Aurora MySQL DB cluster to host the database. Create a DMS replication task to copy the existing data to the target DB cluster. Create a local AWS Schema Conversion Tool (AWS SCT) change data capture (CDC) task to keep the data synchronized. Install the rest of the software on EC2 instances by starting with a compatible base AMI.

D. Deploy an AWS Storage Gateway Volume Gateway on premises. Mount volumes on all on-premises servers. Install the application and the MySQL database on the new volumes. Take regular snapshots. Install all the software on EC2 Instances by starting with a compatible base AMI. Launch a Volume Gateway on an EC2 instance. Restore the volumes from the latest snapshot. Mount the new volumes on the EC2 instances in the case of a failure event.

Suggested Answer: *C*

Community vote distribution

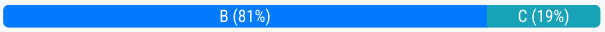

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: B

Tricky one. This is not an on premise migration use case which prompts for answer C. Its a current situation of on premise application which the company wants to continue its state in the requirement of using AWS as DR solution.
https://docs.aws.amazon.com/images/drs/latest/userguide/images/drs-failback-arc.png
https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html
upvoted 27 times

God_Is_Love 1 year, 9 months ago

Moreover, B has least operational over head of just initiating DR solution with replicating agents. C has operational overhead with DMS , SCT ,CDC,migration etc
upvoted 4 times

swadeey 1 year ago

I also agreed with the answer but then see this "The application runs on physical servers that also run other applications. The on-premises application that the company is planning to migrate uses a MySQL database as a data store" just database and physical server has other applications which not mentioned. Also from DR the statement gets changed to Migrate
upvoted 3 times

Untamables Highly Voted 1 year, 11 months ago

Selected Answer: B

https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html
https://docs.aws.amazon.com/drs/latest/userguide/recovery-workflow-gs.html
Option C is wrong. That just mentions the migration method. I think this question asks us the DR architecture between on-premises and AWS cloud.
upvoted 7 times

jimee11 Most Recent 1 month, 2 weeks ago

Selected Answer: B

Poorly worded question. AWS DMS makes sense if we were just migrating the Mysql Database. Since we are migrating the database AND application - we use AWS Replication Agent to migrate the entire server.
upvoted 1 times

altonh 5 months, 2 weeks ago

Selected Answer: C

The answer should be C.

Take note of the statement, "The application runs on physical servers that also run other applications". If you use the Application Migration Service, you will migrate these other applications, which have nothing to do with the application you are trying to protect.

upvoted 2 times

ninomfr64 11 months, 1 week ago

Selected Answer: B

A = to use AWS DRS you first need to set it up in each AWS Region in which you want to use it. installing AWS Replication agent is not enough
B = correct (to me the sentence "Frequently perform failover and fallback from the most recent point in time" is ambiguous as this points to actual failover/failback and not to drills)
C = SCT is not needed wwith same engine db migration. also, install the rest of the software is not enough for app DR
D = Volume Gateway can be used in a Back and Restore DR scenario, but the option D is very confused. Anyway, Storage Gateway for DR requires more overhead with respect to AWS DRS

upvoted 4 times

career360guru 1 year ago

Selected Answer: B

Option B is right option.
Option C only addresses DB instance replication and DR, it does not meet requirements of replicating other applications running on on-premise.

upvoted 1 times

swadeey 1 year ago

Selected answer C changed from B

The application runs on physical servers that also run other applications. The on-premises application that the company is planning to migrate uses a MySQL database as a data store.

upvoted 1 times

severlight 1 year, 1 month ago

Selected Answer: B

Elastic Disaster Recovery does the job

upvoted 1 times

AMohanty 1 year, 3 months ago

C
We are looking for a Business Continuity Solution
Meaning RTO should be low

upvoted 1 times

chikorita 1 year, 3 months ago

but how is failover happening
the very own purpose of DR is its automatic failover which is supported by option B

upvoted 1 times

cmoreira 1 year, 3 months ago

Selected Answer: B

Answer is B.
Questions mentions "least operational overhead" (efforts in the future), and B mentions "Frequently performing...".
However, that is the best-practice for AWS DR (as misleading as it sounds):
https://docs.aws.amazon.com/drs/latest/userguide/failback-overview.html

upvoted 1 times

Gabehcoud 1 year, 4 months ago

Selected Answer: B

the question is a bit misleading, first part says "company is planning for business continuity" the later part of the sentence says "applications are migrating".
nevertheless, we should focus on the word business continuity. Going by that "no migration" is required so choose B.

that is my analysis.

upvoted 3 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

B for BC

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: B

keyword = AWS Elastic Disaster Recovery

B

upvoted 1 times

rbm2023 1 year, 7 months ago

Selected Answer: B

The company is looking for a disaster recovery solution and not a full migration to cloud. In my view the answer should use Elastic Disaster Recovery and not DMS.

References

https://www.cloudthat.com/resources/blog/scalable-cost-effective-cloud-disaster-recovery-with-aws-drs-elastic-disaster-recovery

https://catalog.us-east-1.prod.workshops.aws/workshops/080af3a5-623d-4147-934d-c8d17daba346/en-US/introduction

https://docs.aws.amazon.com/pt_br/mgn/latest/ug/Network-Settings-Video.html

upvoted 2 times

OCHT 1 year, 8 months ago

Selected Answer: C

it appears that option C has the least operational overhead since it involves creating AWS DMS replication servers and a target Amazon Aurora MySQL DB cluster to host the database, creating a DMS replication task to copy existing data to the target DB cluster, creating a local AWS SCT CDC task to keep data synchronized, and installing the rest of the software on EC2 instances by starting with a compatible base AMI. The other options involve additional steps such as setting up replication for all servers (option A), initializing AWS Elastic Disaster Recovery and frequently performing failover and fallbacks (option B), or deploying an AWS Storage Gateway Volume Gateway and mounting volumes on all on-premises servers (option D).

upvoted 3 times

dev112233xx 1 year, 8 months ago

Selected Answer: C

C seems correct to me (DMS with SCT and CDC)

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: B

B has less operational overhead.

upvoted 3 times

Question #149 *Topic 1*

A company is subject to regulatory audits of its financial information. External auditors who use a single AWS account need access to the company's AWS account. A solutions architect must provide the auditors with secure, read-only access to the company's AWS account. The solution must comply with AWS security best practices.

Which solution will meet these requirements?

A. In the company's AWS account, create resource policies for all resources in the account to grant access to the auditors' AWS account. Assign a unique external ID to the resource policy.

B. In the company's AWS account, create an IAM role that trusts the auditors' AWS account. Create an IAM policy that has the required permissions. Attach the policy to the role. Assign a unique external ID to the role's trust policy.

C. In the company's AWS account, create an IAM user. Attach the required IAM policies to the IAM user. Create API access keys for the IAM user. Share the access keys with the auditors.

D. In the company's AWS account, create an IAM group that has the required permissions. Create an IAM user in the company's account for each auditor. Add the IAM users to the IAM group.

Suggested Answer: *B*

Community vote distribution

B (100%)

tatdatpham Highly Voted 1 year, 11 months ago

Selected Answer: B

Option B is the best solution. This solution creates an IAM role that trusts the auditors' AWS account and attaches the required IAM policies to the role. This ensures that the auditors have read-only access to the company's AWS account while ensuring that the company's AWS account is secure and complies with AWS security best practices. Additionally, the unique external ID assigned to the role's trust policy adds an extra layer of security.

upvoted 7 times

duriselvan Most Recent 10 months, 1 week ago

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user.html

upvoted 1 times

duriselvan 10 months, 1 week ago

To create an IAM role that trusts the auditors' AWS account, you can do the following:
Sign in to the AWS Management Console and open the IAM console.
In the navigation pane, choose Roles, and then choose Create role.
Choose the Custom trust policy role type.
In the Custom trust policy section, enter or paste the following trust policy:

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Principal": {
"AWS": "arn:aws:iam::<auditor-account-id>:root"
},
"Action": "sts:AssumeRole"
}
]
}
```

upvoted 1 times

career360guru 1 year ago

Selected Answer: B

Option B

upvoted 1 times

dkcloudguru 1 year, 3 months ago

B is correct

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

its a b

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: B

In the company's AWS account, create an IAM role that trusts the auditors' AWS account.

upvoted 3 times

zozza2023 1 year, 11 months ago

Selected Answer: B

B seems to be the right answer

upvoted 3 times

masetromain 1 year, 11 months ago

Selected Answer: B

The correct answer is B. In the company's AWS account, create an IAM role that trusts the auditors' AWS account. Create an IAM policy that has the required permissions. Attach the policy to the role. Assign a unique external ID to the role's trust policy.

This solution meets the requirement of providing the external auditors with secure, read-only access to the company's AWS account while also complying with AWS security best practices. In this solution, an IAM role is created that trusts the auditors' AWS account and has an IAM policy with the required permissions attached to it. The role's trust policy should include a unique external ID for added security. This allows the external auditors to assume the role and access the resources with the permissions specified in the policy, without the need to share access keys or create individual IAM users for each auditor.

upvoted 3 times

masetromain 1 year, 11 months ago

Option A is incorrect because it grants access to all resources in the company's AWS account and does not provide a way to restrict the permissions that the external auditors have.

Option C is incorrect because it creates an IAM user in the company's account and shares the API access keys with the external auditors, which is not secure and does not comply with AWS security best practices.

Option D is incorrect because it creates an IAM user in the company's account for each auditor, which would be tedious and difficult to manage for the company. It would be more secure and efficient to use an IAM role that trusts the auditors' AWS account instead of creating individual users for each auditor.

upvoted 2 times

zhangyu20000 1 year, 11 months ago

B is correct

upvoted 2 times

Question #150 *Topic 1*

A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.

Which solution will meet these requirements with the LEAST latency?

A. Create a two-node DynamoDB Accelerator (DAX) cluster. Configure an application to read and write data by using DAX.

B. Create a three-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.

C. Create a three-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.

D. Create a single-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.

Suggested Answer: *A*

Community vote distribution

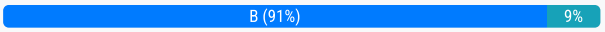

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: B

3 nodes are required for a DAX cluster to be fault-tolerant.
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html
upvoted 19 times

Ganshank Highly Voted 1 year, 10 months ago

This is a poorly framed question with very little attention to how applications are architected in real life. Here's my reasoning:
This being a trading platform, you have a high volume of writes and reads, and stale data is essentially worse than useless. This automatically eliminates all but A, because of the way DAX performs. DAX caches data from the first query, and subsequent queries will continue to receive that cached data regardless of whether it has been updated in DynamoDB. This behavior continues till cache eviction. The only way around it is to read and write data using DAX.
Here's the curveball - the solution must be HA, which eliminates A and D, leaving only B & C. And between B & C, you really want to use DAX for reading and DynamoDB for writing. So final answer is B - if you want to get certified.
Applying this solution in real world however will cause you a lot of pain and grief!
upvoted 13 times

jainparag1 1 year, 7 months ago

Cahing in DAX is always write through. Correct answer is B.
upvoted 2 times

frfavoreto 1 year, 9 months ago

Totally agree.

But an additional issue with the question is the fact that it requires High Availability, not Fault Tolerance. These are quite different concepts and, at least up to this point, there would be no need for 3x DAX instances (in theory).
upvoted 1 times

ThachNguyen Most Recent 9 months ago

Selected Answer: B

B is Correct.
- To achieve high availability for your application, we recommend that you provision your DAX cluster with at least three nodes. Ref: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.consistency.html#DAX.consistency.nodes
- If the request specifies eventually consistent reads (the default behavior), it tries to read the item from DAX.
- With these operations, data is first written to the DynamoDB table, and then to the DAX cluster. The operation is successful only if the data is successfully written to both the table and to DAX.
Ref: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.html#DAX.concepts.request-processing

upvoted 1 times

saggy4 1 year, 4 months ago

Selected Answer: B

DAX is cache and can only be used to read so A and C are out.
Between B and D the question says Highly Available so we will select B (three node) instead of D (single node).

So correct answer B
upvoted 5 times

ninomfr64 1 year, 5 months ago

A = 2 nodes DAX is not fault-tolerant
B = correct (write-around strategy ensure lower latency)
C = write-through strategy can have higher latency
D = 1 node DAX is not fault-tolerant
upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: B

Option B is the best option. Though Option A is also possible solution.
upvoted 1 times

MRamos 1 year, 6 months ago

Selected Answer: B

The breakpoint is latency.

You write throught DAX, but for latency sensitive apps, AWS instruct write directly on DynamoDB instead on DAX.

"For applications that are sensitive to latency, writing through DAX incurs an extra network hop. So a write to DAX is a little slower than a write directly to DynamoDB. If your application is sensitive to write latency, you can reduce the latency by writing directly to DynamoDB instead. For more information, see Write-around."

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.consistency.html#DAX.consistency.strategies-for-writes
upvoted 1 times

amarbeg 1 year, 7 months ago

Option A would be the least latency solution for this use case. Using a two node DAX cluster with the application reading and writing via DAX provides:

Caching of both reads and writes within the DAX cluster nodes. This eliminates the need to go directly to DynamoDB for reads and writes, reducing latency.

Redundancy with two nodes to ensure high availability of the cache.

The other options would lead to some reads or writes still going directly to DynamoDB rather than being fully served from the lower latency cached data in DAX. This could increase latency compared to option A. A single node DAX cluster would work but lacks the redundancy needed for high availability.

DAX is fully managed, in-memory cache for DynamoDB that delivers low-latency data access. By caching the entire dataset in-memory across nodes, it can serve requests much faster than going to the DynamoDB tables on every request. The AWS documentation provides more details on how to configure DAX and monitor latency metrics.
upvoted 1 times

covabix879 1 year, 8 months ago

Selected Answer: A

Question only ask for High Availability, not Fault Tolerant. You need 3 nodes only for the latter. You must write through to keep data getting stale as mentioned by Ganshank. I would go with two-node cluster as strong consistency adds extra latency as number of clusters increase. So for this question best answer should be A.
upvoted 1 times

dkcloudguru 1 year, 9 months ago

Option B is correct: DAX is also used for caching so it improves the performance and for production 3 nodes are strongly recommended so i ll go with B.

upvoted 2 times

duriselvan 1 year, 9 months ago

https://aws.amazon.com/blogs/database/amazon-dynamodb-accelerator-dax-a-read-throughwrite-through-cache-for-dynamodb/

upvoted 1 times

duriselvan 1 year, 9 months ago

sorry guys a is wrong ans: B is correct ans Important

For production usage, we strongly recommend using DAX with at least three nodes, where each node is placed in different Availability Zones. Three nodes are required for a DAX cluster to be fault-tolerant.

A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.

upvoted 1 times

duriselvan 1 year, 9 months ago

A is Ans :

Read replicas serve two additional purposes:

Scalability. If you have a large number of application clients that need to access DAX concurrently, you can add more replicas for read-scaling. DAX spreads the load evenly across all the nodes in the cluster. (Another way to increase throughput is to use larger cache node types.)

High availability. In the event of a primary node failure, DAX automatically fails over to a read replica and designates it as the new primary. If a replica node fails, other nodes in the DAX cluster can still serve requests until the failed node can be recovered. For maximum fault tolerance, you should deploy read replicas in separate Availability Zones. This configuration ensures that your DAX cluster can continue to function, even if an entire Availability Zone becomes unavailable.

upvoted 1 times

AMohanty 1 year, 9 months ago

A

Once u enable DAX you cant directly write onto or Read from Dynamo DB.

upvoted 2 times

ggrodskiy 1 year, 11 months ago

Correct B.

upvoted 1 times

Just_Ninja 1 year, 11 months ago

Selected Answer: B

AWS recommend 3 nodes for production workloads.

So it must B

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

B for DAX HA

upvoted 1 times

Question #151 *Topic 1*

A company has migrated an application from on premises to AWS. The application frontend is a static website that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). The application backend is a Python application that runs on three EC2 instances behind another ALB. The EC2 instances are large, general purpose On-Demand Instances that were sized to meet the on-premises specifications for peak usage of the application.

The application averages hundreds of thousands of requests each month. However, the application is used mainly during lunchtime and receives minimal traffic during the rest of the day.

A solutions architect needs to optimize the infrastructure cost of the application without negatively affecting the application availability.

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

A. Change all the EC2 instances to compute optimized instances that have the same number of cores as the existing EC2 instances.

B. Move the application frontend to a static website that is hosted on Amazon S3.

C. Deploy the application frontend by using AWS Elastic Beanstalk. Use the same instance type for the nodes.

D. Change all the backend EC2 instances to Spot Instances.

E. Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances.

Suggested Answer: *BE*

Community vote distribution

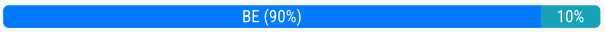

severlight Highly Voted 1 year, 7 months ago

Selected Answer: BE

Burstable instances let you save costs, you pay for some baseline - say 40 percent, if the instance is utilized less - credits get accumulated. So, it is good for workloads with changing CPU loads.

upvoted 10 times

kiran15789 Highly Voted 2 years, 4 months ago

Selected Answer: BE

Burstable EC2 instances, also known as T instances, provide a baseline level of CPU performance with the ability to burst CPU usage when additional cycles are available. They are designed for workloads that do not require sustained high CPU performance but occasionally need more CPU power. Burstable instances can be a cost-effective option for workloads that have moderate CPU requirements but still require flexibility to handle occasional spikes in demand.

upvoted 5 times

sse69 Most Recent 1 year, 1 month ago

Uhm, S3 static website with a Python backend? Am I missing something? How can S3 interact with a backend?

upvoted 1 times

helloworldabc 10 months, 1 week ago

just B,E

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: BE

Option B and E

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: BE

it's BE

upvoted 3 times

rbm2023 2 years, 1 month ago

Selected Answer: BE

You cannot move all backend to Spot Instances this will break the requirement for not affecting the application availability.
You can improve by moving the static site to S3, front end, and change the on demand instances to burst capacity.

upvoted 4 times

OCHT 2 years, 2 months ago

Selected Answer: BE

Amazon EC2 Spot Instances allow you to take advantage of unused EC2 capacity in the AWS Cloud at a steep discount compared to On-Demand Instance prices. Spot Instances are well-suited for workloads that can be interrupted, such as batch processing, data analysis, and image or video processing. They can also be used for fault-tolerant workloads that can withstand the loss of an instance, such as web services or stateless applications.

upvoted 2 times

OCHT 2 years, 2 months ago

Option C suggests deploying the application frontend using AWS Elastic Beanstalk and using the same instance type for the nodes. Elastic Beanstalk is a fully managed service that makes it easy to deploy, run, and scale applications. It automatically handles the deployment and management of the underlying infrastructure, including capacity provisioning, load balancing, and auto-scaling. However, using Elastic Beanstalk with the same instance type as the existing EC2 instances may not necessarily reduce costs.

upvoted 1 times

OCHT 2 years, 2 months ago

Option E suggests deploying the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances. Burstable instances provide a baseline level of CPU performance with the ability to burst above the baseline when needed. This can be a cost-effective option for workloads that have variable CPU usage and can benefit from the ability to burst during periods of high demand. However, if the workload consistently requires high CPU usage, using burstable instances may not provide significant cost savings compared to using larger general purpose instances.

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: BE

BE makes the most sense here

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: BE

Burstable because peak performance is needed at lunch time and its cost effective based on this -
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/burstable-performance-instances.html
S3 static website hosting is cost effective

upvoted 5 times

tatdatpham 2 years, 4 months ago

Selected Answer: BE

The correct answer is B, E.
Option B of moving the frontend to a static website hosted on Amazon S3 will reduce the cost of running the frontend, as S3 is a lower cost storage option than EC2 instances.
Option E of deploying the backend Python application to general purpose burstable EC2 instances will ensure that the backend EC2 instances have the capacity to handle spikes in usage, as burstable instances are designed to handle unpredictable workloads. This will help to optimize the cost of running the backend, as burstable instances are less expensive than On-Demand instances and more cost-effective than Spot instances.

upvoted 1 times

Untamables 2 years, 5 months ago

Selected Answer: BE

B and E.
Option D is wrong. A spot instance is not appropriate for a production server.
By the way, I would like another option that mentions changing the backend Python API Gateway and Lambda because Option B mentions changing the frontend serverless. I think this question is a typical use case of the serverless architecture.

upvoted 4 times

vsk12 2 years, 5 months ago

Selected Answer: BE

Correct answers are
B & E

Option B as S3 is a cost-effective storage solution for static websites.
Option E as burstable general-purpose instances provides a cost-effective solution for this kind of workload.
upvoted 2 times

masetromain 2 years, 5 months ago

Selected Answer: BD

B. Move the application frontend to a static website that is hosted on Amazon S3.
D. Change all the backend EC2 instances to Spot Instances.

Step 1: Moving the application frontend to a static website that is hosted on Amazon S3 will reduce the cost and increase the scalability of the application. S3 is a highly scalable object storage service that can handle large amounts of data and traffic at a lower cost than running EC2 instances.

Step 2: Changing the backend EC2 instances to Spot Instances can help reduce cost without negatively affecting the application availability. Spot Instances allow customers to bid on unused Amazon EC2 capacity, which can result in significant cost savings. You can also use AWS Auto Scaling to automatically increase or decrease the number of Spot Instances based on the application's traffic.
upvoted 4 times

masetromain 2 years, 5 months ago

Option A, C: Changing to compute optimized instances or using Elastic Beanstalk will not help reducing the cost, it will only change the instances type and not helping the cost optimization.
Option E: Deploying the backend Python application to general purpose burstable EC2 instances will not help reducing the cost, as it still using On-Demand instances.

It is important to note that using spot instances comes with the risk of instances being terminated when the spot price goes up. To mitigate this risk, you could use the EC2 Auto Scaling group with a combination of on-demand and spot instances. This way, if a spot instance is terminated, the Auto Scaling group can automatically replace it with an on-demand instance to ensure the application is always available.
upvoted 1 times

zhangyu20000 2 years, 5 months ago

BE are correct
A: Compute optimized instance is expensive than burstable instance
B: S3 hosted static web server is cheaper
C: Not save money
D: Spot instance affect availibility
E: Burstable EC2 is cheaper
upvoted 2 times

masetromain 2 years, 5 months ago

To mitigate this risk, you could use the EC2 Auto Scaling group with a combination of on-demand and spot instances. This way, if a spot instance is terminated, the Auto Scaling group can automatically replace it with an on-demand instance to ensure the application is always available.
upvoted 1 times

Question #152 *Topic 1*

A company is running an event ticketing platform on AWS and wants to optimize the platform's cost-effectiveness. The platform is deployed on Amazon Elastic Kubernetes Service (Amazon EKS) with Amazon EC2 and is backed by an Amazon RDS for MySQL DB instance. The company is developing new application features to run on Amazon EKS with AWS Fargate.

The platform experiences infrequent high peaks in demand. The surges in demand depend on event dates.

Which solution will provide the MOST cost-effective setup for the platform?

A. Purchase Standard Reserved Instances for the EC2 instances that the EKS cluster uses in its baseline load. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet predicted peak load for the year.

B. Purchase Compute Savings Plans for the predicted medium load of the EKS cluster. Scale the cluster with On-Demand Capacity Reservations based on event dates for peaks. Purchase 1-year No Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale out database read replicas during peaks.

C. Purchase EC2 Instance Savings Plans for the predicted base load of the EKS cluster. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale up the DB instance manually during peaks.

D. Purchase Compute Savings Plans for the predicted base load of the EKS cluster. Scale the cluster with Spot Instances to handle peaks. Purchase 1-year All Upfront Reserved Instances for the database to meet the predicted base load. Temporarily scale up the DB instance manually during peaks.

Suggested Answer: *B*

Community vote distribution

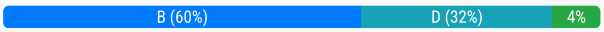

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: B

Option A, C and D are wrong. They all mention using spot instances and EKS based on EC2. A spot instance is not appropriate for a production server and the company is developing new application designed for AWS Fargate, which means we must plan the future cost improvement including AWS Fargate.

https://aws.amazon.com/savingsplans/compute-pricing/

upvoted 20 times

zhangyu20000 Highly Voted 2 years, 5 months ago

B is correct. Compute saving plan will also cover Fargate

A: use spot instance is not reliable

CD: manually scale up DB

upvoted 12 times

bhanus Most Recent 6 months, 1 week ago

Selected Answer: D

Compute Savings Plans for EKS Base Load:

Spot Instances for Peaks:

1-Year All Upfront Reserved Instances for Database Base Load:

upvoted 1 times

nelgeozcin 7 months, 2 weeks ago

B - Fargate cannot support Spot - https://docs.aws.amazon.com/eks/latest/userguide/fargate.html

upvoted 1 times

Sin_Dan 8 months, 2 weeks ago

Selected Answer: D

No brainer, it's D. C doesn't provide any cost-effectiveness!

upvoted 4 times

FZA24 8 months, 2 weeks ago

Selected Answer: B

A wrong : Spot Instances to handle peaks
B: correct
C & D wrong : Temporarily scale up the DB instance manually during peaks.
upvoted 1 times

vip2 11 months, 1 week ago

Selected Answer: D

D looks more correctable
Mainly diff. between B and D is
predicable workload-- all upfront
no specific for read-replication traffic
On-Demand Capacity Reservations ensure availability during peak times without long-term commitments. but no cost-effective
upvoted 4 times

helloworldabc 10 months, 1 week ago

just B
upvoted 1 times

thotwielder 1 year, 1 month ago

Selected Answer: D

It's between b and d. D is more cost effective because of spot instances. And B is wrong because there is no reason to scale read replicas for RDS (the question doesn't say read only load)
upvoted 6 times

Dgix 1 year, 3 months ago

Selected Answer: D

I really don't understand why people are saying that Spot instances aren't suitable for production. There is a two-minute respite before they shut down, and since the application is not said to be stateful, this is plenty of time for a single request/response cycle.

With this in mind, the correct solution is D.
upvoted 6 times

Keval12345 1 year, 2 months ago

slightly difference betwee B and D {other than spot instances ofcourse}. Since the platform experinces peaks, might be a better idea to go for savings plan with medium load
upvoted 2 times

saggy4 1 year, 4 months ago

Selected Answer: B

A and C: The company will have a mix of EKS on EC2 and EKS Fargate hence reserved instance is not possible as it will cover only EKS on EC2 hence A and C are out

Between B and C:
C seems to save the most cost, but during peak load spot instances (both EC2 or Fargate) will not provide guaranteed availability. Hence we should go ahead with B.

Correct Answer: B
upvoted 2 times

AWSLord32 1 year, 4 months ago

Selected Answer: C

C is the right answer. Everything about B is wrong: Compute savings plan is more expensive than RI, on demand more expensive than spot for peaks and no upfront more exponsive than all upfront.
upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: D

The scenario ask for the most cost-effective setup. Thus:
A = RI doesn't cover Fargate
B = ODCR doesn't bring cost benefits, they just ensure you have capacity. Read replicas are for read only, I would expect workload peaks includes writes so this is not saving money nor fully helping with capacity needs
C = EC2 Saving Plans do not cover Fargate
D = correct (this is the most cost-effective setup, Compute Savings Plans apply to both EC2 and Fargate, Spot Instances applies to both EC2 and

Fargate, All Upfront Reserved Instances is most cost effective option for RDS. Manually scaling RDS adds a lot of overhead, but this is not the point of the question)

upvoted 6 times

ninomfr64 1 year, 5 months ago

Also, for a temporarily limited change it is easier to manually vertically scale your instance rather than adding Read replicas as adding replicas to a single instance requires to change your app to send read requests to the reader endpoint and not to the cluster (aka writer) endpoint

upvoted 2 times

Jay_2pt0_1 1 year, 6 months ago

I might be leaning toward D as it does ask for the. most cost-effective solution

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: D

Compute saving plans are more cost effective so B or D are right two options.
Between B and D - Spot instances offers better cost and Fargate supports spot instances
https://aws.amazon.com/blogs/aws/aws-fargate-spot-now-generally-available/
Option B says, scale RDS Read-Replica for based on events which may not work as workload description does not mentioned that peak load is only read traffic. So D is best and most cost effective solution.

upvoted 8 times

Hyperdanny 1 year, 11 months ago

Selected Answer: C

I am leaning towards C, since Instance savings provide the biggest discount.
I also couldn't find a way to scale EKS based on dates, which B suggests: "Scale the cluster with On-Demand Capacity Reservations based on event dates for peaks"

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

its a b

upvoted 2 times

SkyZeroZx 2 years ago

Selected Answer: B

Option A, C and D are wrong. They all mention using spot instances and EKS based on EC2. A spot instance is not appropriate for a production server and the company is developing new application designed for AWS Fargate, which means we must plan the future cost improvement including AWS Fargate.

upvoted 4 times

Question #153 *Topic 1*

A company has deployed an application on AWS Elastic Beanstalk. The application uses Amazon Aurora for the database layer. An Amazon CloudFront distribution serves web requests and includes the Elastic Beanstalk domain name as the origin server. The distribution is configured with an alternate domain name that visitors use when they access the application.

Each week, the company takes the application out of service for routine maintenance. During the time that the application is unavailable, the company wants visitors to receive an informational message instead of a CloudFront error message.

A solutions architect creates an Amazon S3 bucket as the first step in the process.

Which combination of steps should the solutions architect take next to meet the requirements? (Choose three.)

A. Upload static informational content to the S3 bucket.

B. Create a new CloudFront distribution. Set the S3 bucket as the origin.

C. Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI).

D. During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete.

E. During the weekly maintenance, create a cache behavior for the S3 origin on the new distribution. Set the path pattern to \ Set the precedence to 0. Delete the cache behavior when the maintenance is complete.

F. During the weekly maintenance, configure Elastic Beanstalk to serve traffic from the S3 bucket.

Suggested Answer: *ACD*

Community vote distribution

ACD (100%)

masetromain Highly Voted 2 years, 5 months ago

Selected Answer: ACD

A. Upload static informational content to the S3 bucket.
C. Set the S3 bucket as a second origin in the original CloudFront distribution. Configure the distribution and the S3 bucket to use an origin access identity (OAI).
D. During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete.

Step 1: The solutions architect should upload static informational content to the S3 bucket, this content will be shown to the users when the application is down for maintenance.

Step 2: The solutions architect should set the S3 bucket as a second origin in the original CloudFront distribution. To keep the S3 bucket secure, the solutions architect should configure the distribution and the S3 bucket to use an origin access identity (OAI). This will ensure that only CloudFront has access to the S3 bucket.

upvoted 16 times

masetromain 2 years, 5 months ago

Step 3: During the weekly maintenance, the solutions architect should edit the default cache behavior of the CloudFront distribution to use the S3 origin. This will redirect all incoming traffic to the S3 bucket and show the static informational content to the users. Once the maintenance is complete, the solutions architect should revert the change back to the original Elastic Beanstalk origin.

Option B: Creating a new CloudFront distribution and setting the S3 bucket as the origin is unnecessary and could cause confusion for the users.
Option E: During the weekly maintenance, creating a cache behavior for the S3 origin on the new distribution is unnecessary, it is more complex and prone to human error.
Option F: Configuring Elastic Beanstalk to serve traffic from the S3 bucket is not necessary because CloudFront is already being used as the web request server.

upvoted 5 times

carpa_jo Most Recent 1 year, 6 months ago

Selected Answer: ACD

From the given options ACD makes the most sense.
In real life the CloudFront feature to show custom error responses might make a lot more sense:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html#custom-error-pages-procedure
This would avoid the manual steps and by that is less prone to human errors.
upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: ACD

A, C and D is correct.
upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: ACD

CacheBehaviour defines path and origin
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: ACD

ACD morelikely
upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: ACD

A C D
E is good option but is more overhead and propone error human then C is more accesible
upvoted 2 times

Jesuisleon 2 years ago

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: ACD

ACD is the best fit
upvoted 3 times

Musk 2 years, 4 months ago

Selected Answer: ACD

About E, the lowest possible value for the "Origin Priority" field in AWS CloudFront is 1
upvoted 4 times

sam2ng 8 months, 2 weeks ago

behavior precedence can be set to zero
upvoted 1 times

zozza2023 2 years, 5 months ago

Selected Answer: ACD

ACD is correct
upvoted 4 times

zhangyu20000 2 years, 5 months ago

ABD is correct
upvoted 1 times

zhangyu20000 2 years, 5 months ago

ACD is correct
upvoted 2 times

Question #154 *Topic 1*

A company gives users the ability to upload images from a custom application. The upload process invokes an AWS Lambda function that processes and stores the image in an Amazon S3 bucket. The application invokes the Lambda function by using a specific function version ARN.

The Lambda function accepts image processing parameters by using environment variables. The company often adjusts the environment variables of the Lambda function to achieve optimal image processing output. The company tests different parameters and publishes a new function version with the updated environment variables after validating results. This update process also requires frequent changes to the custom application to invoke the new function version ARN. These changes cause interruptions for users.

A solutions architect needs to simplify this process to minimize disruption to users.

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

A. Directly modify the environment variables of the published Lambda function version. Use the SLATEST version to test image processing parameters.

B. Create an Amazon DynamoDB table to store the image processing parameters. Modify the Lambda function to retrieve the image processing parameters from the DynamoDB table.

C. Directly code the image processing parameters within the Lambda function and remove the environment variables. Publish a new function version when the company updates the parameters.

D. Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Suggested Answer: *D*

Community vote distribution

D (100%)

tatdatpham Highly Voted 1 year, 11 months ago

Selected Answer: D

D is correct

By using a function alias, the custom application invokes the latest version of the Lambda function without the need to modify the application code every time the company updates the image processing parameters. This reduces the risk of causing interruptions for users.

upvoted 12 times

masetromain Highly Voted 1 year, 11 months ago

Selected Answer: D

D. Create a Lambda function alias. Modify the client application to use the function alias ARN. Reconfigure the Lambda alias to point to new versions of the function when the company finishes testing.

Creating a Lambda function alias allows the solutions architect to change the version of the Lambda function that the alias points to without modifying the client application. This eliminates the need for frequent updates to the custom application and minimizes disruption to users. The solutions architect can test different parameters by using different versions of the function and reconfigure the alias to point to the new version after validating results. This allows the company to update the image processing parameters without affecting the users.

upvoted 5 times

masetromain 1 year, 11 months ago

Option A: Directly modifying the environment variables of the published Lambda function version would cause all clients to use the updated environment variables immediately and would not allow for testing.

Option B: Using DynamoDB to store image processing parameters increases complexity and operational overhead, and it would not eliminate the need for updating the custom application.

Option C: Directly coding the image processing parameters within the Lambda function and publishing new versions would not eliminate the need for updating the custom application.

upvoted 2 times

career360guru Most Recent 1 year ago

Selected Answer: D

Option D has least operational overhead.

upvoted 1 times

edder 1 year ago

Selected Answer: D

https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html

upvoted 1 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: D

Look for ALIAS

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: D

D

B is ok, but more overhead

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: D

keyword = Lambda ALIAS

then D

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: D

Create a Lambda function alias.

upvoted 1 times

zhangyu20000 1 year, 11 months ago

D is correct

upvoted 1 times

Question #155 *Topic 1*

A global media company is planning a multi-Region deployment of an application. Amazon DynamoDB global tables will back the deployment to keep the user experience consistent across the two continents where users are concentrated. Each deployment will have a public Application Load Balancer (ALB). The company manages public DNS internally. The company wants to make the application available through an apex domain.

Which solution will meet these requirements with the LEAST effort?

A. Migrate public DNS to Amazon Route 53. Create CNAME records for the apex domain to point to the ALB. Use a geolocation routing policy to route traffic based on user location.

B. Place a Network Load Balancer (NLB) in front of the ALMigrate public DNS to Amazon Route 53. Create a CNAME record for the apex domain to point to the NLB's static IP address. Use a geolocation routing policy to route traffic based on user location.

C. Create an AWS Global Accelerator accelerator with multiple endpoint groups that target endpoints in appropriate AWS Regions. Use the accelerator's static IP address to create a record in public DNS for the apex domain.

D. Create an Amazon API Gateway API that is backed by AWS Lambda in one of the AWS Regions. Configure a Lambda function to route traffic to application deployments by using the round robin method. Create CNAME records for the apex domain to point to the API's URL.

Suggested Answer: *C*

Community vote distribution

C (100%)

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: C

No, an apex domain cannot use CNAME records in AWS. This is because of the way DNS resolution works. A CNAME record specifies an alias for a domain name, which points to the canonical name of another domain. However, the DNS standard does not allow CNAME records for apex domains, as they should only have A or AAAA records.

When you try to create a CNAME record for an apex domain in AWS Route 53, you will receive an error message indicating that the record set type is not valid for the apex domain. To work around this limitation, you can use an alias record instead.

upvoted 22 times

zhangyu20000 Highly Voted 1 year, 11 months ago

C is correct

ABD all have CNAME record that is not allowed for apex domain

upvoted 10 times

career360guru Most Recent 1 year ago

Selected Answer: C

Option C

upvoted 2 times

yuliaqwerty 1 year ago

C https://aws.amazon.com/blogs/networking-and-content-delivery/solving-dns-zone-apex-challenges-with-third-party-dns-providers-using-aws/

upvoted 3 times

[Removed] 1 year, 2 months ago

You can create alias record for apex domain in route 53 However the question is asking about least effort and the client is managing domain internally

upvoted 1 times

Explorer_30 1 year, 3 months ago

The answer is C

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

C
no CNAME for apex
upvoted 2 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: C

A , B no seems because reference geolocation
D no seems because apex domain with API Gateway ?
then C Global Accelerator is good option
upvoted 1 times

chikorita 1 year, 6 months ago

fun fact: CNAME records does not support APEX domain
which simply rules out the options with CNAME in it
answer is C
upvoted 4 times

mfsec 1 year, 9 months ago

Selected Answer: C

Create an AWS Global Accelerator accelerator with multiple endpoint groups that target endpoints in appropriate AWS Regions.
upvoted 3 times

masetromain 1 year, 11 months ago

Selected Answer: C

C. Create an AWS Global Accelerator accelerator with multiple endpoint groups that target endpoints in appropriate AWS Regions. Use the accelerator's static IP address to create a record in public DNS for the apex domain.

This solution meets the requirements with the least effort because it uses AWS Global Accelerator, which automatically routes traffic to the optimal endpoint based on health and geography, eliminating the need for manual configuration or additional routing policies. It also eliminates the need to create a CNAME record for the apex domain to point to the ALB or NLB's IP address, which can be less efficient and less reliable.
upvoted 5 times

masetromain 1 year, 11 months ago

A. Migrate public DNS to Amazon Route 53. Create CNAME records for the apex domain to point to the ALB. Use a geolocation routing policy to route traffic based on user location.
While this solution uses Route 53 and geolocation routing, it requires manual configuration and maintenance of the routing policy and could introduce additional latency as traffic is routed through the ALB first.

B. Place a Network Load Balancer (NLB) in front of the ALB. Migrate public DNS to Amazon Route 53. Create a CNAME record for the apex domain to point to the NLB's static IP address. Use a geolocation routing policy to route traffic based on user location.
This solution is similar to the first one, but it uses a Network Load Balancer (NLB) instead of an Application Load Balancer (ALB). It has the same downsides as the first solution.
upvoted 1 times

masetromain 1 year, 11 months ago

D. Create an Amazon API Gateway API that is backed by AWS Lambda in one of the AWS Regions. Configure a Lambda function to route traffic to application deployments by using the round robin method. Create CNAME records for the apex domain to point to the API's URL.

This solution uses Amazon API Gateway and AWS Lambda to route traffic, but the round-robin method is not the best way to ensure optimal performance and availability for a multi-region deployment. Additionally, routing traffic through a Lambda function can introduce additional latency.

AWS Global Accelerator is a more efficient solution that automatically routes traffic to the optimal endpoint based on health and geography, eliminating the need for manual configuration or additional routing policies.
upvoted 1 times

Question #156 *Topic 1*

A company is developing a new serverless API by using Amazon API Gateway and AWS Lambda. The company integrated the Lambda functions with API Gateway to use several shared libraries and custom classes.

A solutions architect needs to simplify the deployment of the solution and optimize for code reuse.

Which solution will meet these requirements?

A. Deploy the shared libraries and custom classes into a Docker image. Store the image in an S3 bucket. Create a Lambda layer that uses the Docker image as the source. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the Lambda layer.

B. Deploy the shared libraries and custom classes to a Docker image. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Create a Lambda layer that uses the Docker image as the source. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the Lambda layer.

C. Deploy the shared libraries and custom classes to a Docker container in Amazon Elastic Container Service (Amazon ECS) by using the AWS Fargate launch type. Deploy the API's Lambda functions as Zip packages. Configure the packages to use the deployed container as a Lambda layer.

D. Deploy the shared libraries, custom classes, and code for the API's Lambda functions to a Docker image. Upload the image to Amazon Elastic Container Registry (Amazon ECR). Configure the API's Lambda functions to use the Docker image as the deployment package.

Suggested Answer: *B*

Community vote distribution

lunt Highly Voted 2 years, 4 months ago

Selected Answer: D

Don't understand why so many people are choosing B. Read up. A container image cannot be used with Lambda layers. That means A B C are out instantly. Its literally one of the first things they mention about Lamba layers. Answer is D and ABC simply impossible to configure.

https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html

upvoted 45 times

titi_r 1 year, 1 month ago

You can create a Lambda function from an ECR image, but you CANNOT create a Lambda function layer from an ECR image!

upvoted 4 times

Gabehcoud 1 year, 10 months ago

https://aws.amazon.com/blogs/compute/working-with-lambda-layers-and-extensions-in-container-images/

Previously, Lambda functions were packaged only as .zip archives. This includes functions created in the AWS Management Console. You can now also package and deploy Lambda functions as container images.

You can use familiar container tooling such as the Docker CLI with a Dockerfile to build, test, and tag images locally. Lambda functions built using container images can be up to 10 GB in size. You push images to an Amazon Elastic Container Registry (ECR) repository, a managed AWS container image registry service. You create your Lambda function, specifying the source code as the ECR image URL from the registry.

upvoted 3 times

rtgfdv3 2 years, 3 months ago

https://aws.amazon.com/blogs/compute/working-with-lambda-layers-and-extensions-in-container-images/

upvoted 3 times

c73bf38 2 years, 4 months ago

B suggests deploying the shared libraries and custom classes to a Docker image, uploading it to Amazon Elastic Container Registry (Amazon ECR), creating a Lambda layer that uses the Docker image as the source, and deploying the API's Lambda functions as Zip packages. Configuring the packages to use the Lambda layer simplifies deployment, and the Docker image allows for code reuse. This option takes advantage of the built-in features provided by AWS API Gateway and Lambda, making it the optimal solution.

upvoted 5 times

c73bf38 2 years, 4 months ago

The requirement is code reuse:
https://aws.amazon.com/blogs/compute/working-with-lambda-layers-and-extensions-in-container-images/
Lambda functions packaged as container images do not support adding Lambda layers to the function configuration. However, there are a number of solutions to use the functionality of Lambda layers with container images. You take on the responsible for packaging your preferred runtimes and dependencies as a part of the container image during the build process.

upvoted 4 times

Untamables Highly Voted 2 years, 5 months ago

Selected Answer: D

Option A, B and C are wrong. An AWS Lambda Layer does not support a Docker image or a deployed container as the source.
https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html
https://aws.amazon.com/blogs/compute/working-with-lambda-layers-and-extensions-in-container-images/

upvoted 8 times

albert_kuo Most Recent 3 months, 3 weeks ago

Selected Answer: D

Lambda Layer does not support Docker image.

upvoted 1 times

kgpoj 10 months, 3 weeks ago

Selected Answer: D

If any of you ever really worked in lambda with docker image, you will instantly choose D without hesitation.

zipped package can be deployed straightaway and it doesn't need a container. Don't get those two things(lambda zip deployment vs lambda container deployment) mixed up

upvoted 1 times

zolthar_z 11 months, 1 week ago

Selected Answer: D

Please read the requirement, "simplify the deployment" with D you need only to maintain the docker image, with B you need to maintain the docker image and the process to deploy the lambda as ZIP Packages.

upvoted 1 times

Nicoben 1 year, 6 months ago

Selected Answer: B

Option B is the right one, see: https://docs.aws.amazon.com/lambda/latest/dg/images-create.html

upvoted 2 times

career360guru 1 year, 6 months ago

Selected Answer: D

Option D

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: D

check lunt's answer

upvoted 1 times

rlf 1 year, 8 months ago

B.
* A Lambda layer is a .zip file archive that contains supplementary code or data. Layers usually contain library dependencies, a custom runtime, or configuration files.
* Lambda functions packaged as container images do not support adding Lambda layers to the function configuration.However, there are a number of solutions to use the functionality of Lambda layers with container images. You take on the responsible for packaging your preferred runtimes and dependencies as a part of the container image during the build process.

upvoted 2 times

dkcloudguru 1 year, 9 months ago

Ans is D: https://aws.amazon.com/blogs/compute/working-with-lambda-layers-and-extensions-in-container-images/#:~:text=Lambda%20functions%20packaged%20as%20container,Lambda%20layers%20with%20container%20images.

upvoted 1 times

Gabehcoud 1 year, 10 months ago

Answer B.
Previously, Lambda functions were packaged only as .zip archives. This includes functions created in the AWS Management Console. You can now also package and deploy Lambda functions as container images.

You can use familiar container tooling such as the Docker CLI with a Dockerfile to build, test, and tag images locally. Lambda functions built using container images can be up to 10 GB in size. You push images to an Amazon Elastic Container Registry (ECR) repository, a managed AWS container image registry service. You create your Lambda function, specifying the source code as the ECR image URL from the registry.
upvoted 2 times

vn_thanhtung 1 year, 10 months ago
https://www.youtube.com/watch?v=17R0vN8bt-0
upvoted 1 times

ggrodskiy 1 year, 11 months ago
Correct B.
upvoted 2 times

NikkyDicky 1 year, 12 months ago
Selected Answer: D
D
layers not supported w container-based lambdas
upvoted 1 times

pupsik 2 years ago
Selected Answer: D
Docker images cannot be used in Lambda layers.
upvoted 1 times

Jackhemo 2 years ago
Selected Answer: B
From olabiba.ai: Overall, option B provides a streamlined approach to optimize code reuse by centralizing the shared code in a Docker image and using a Lambda layer to share it across multiple functions.
upvoted 1 times

Roontha 2 years ago
Answer : B
upvoted 1 times

rbm2023 2 years, 1 month ago
Selected Answer: B
"Lambda functions packaged as container images do not support adding Lambda layers to the function configuration. However, there are a number of solutions to use the functionality of Lambda layers with container images. You take on the responsible for packaging your preferred runtimes and dependencies as a part of the container image during the build process."
https://aws.amazon.com/blogs/compute/working-with-lambda-layers-and-extensions-in-container-images/
upvoted 6 times

Question #157 *Topic 1*

A manufacturing company is building an inspection solution for its factory. The company has IP cameras at the end of each assembly line. The company has used Amazon SageMaker to train a machine learning (ML) model to identify common defects from still images.

The company wants to provide local feedback to factory workers when a defect is detected. The company must be able to provide this feedback even if the factory's internet connectivity is down. The company has a local Linux server that hosts an API that provides local feedback to the workers.

How should the company deploy the ML model to meet these requirements?

A. Set up an Amazon Kinesis video stream from each IP camera to AWS. Use Amazon EC2 instances to take still images of the streams. Upload the images to an Amazon S3 bucket. Deploy a SageMaker endpoint with the ML model. Invoke an AWS Lambda function to call the inference endpoint when new images are uploaded. Configure the Lambda function to call the local API when a defect is detected.

B. Deploy AWS IoT Greengrass on the local server. Deploy the ML model to the Greengrass server. Create a Greengrass component to take still images from the cameras and run inference. Configure the component to call the local API when a defect is detected.

C. Order an AWS Snowball device. Deploy a SageMaker endpoint the ML model and an Amazon EC2 instance on the Snowball device. Take still images from the cameras. Run inference from the EC2 instance. Configure the instance to call the local API when a defect is detected.

D. Deploy Amazon Monitron devices on each IP camera. Deploy an Amazon Monitron Gateway on premises. Deploy the ML model to the Amazon Monitron devices. Use Amazon Monitron health state alarms to call the local API from an AWS Lambda function when a defect is detected.

Suggested Answer: *D*

Community vote distribution

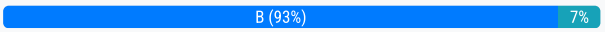

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: B

Offline operation: AWS IoT Greengrass supports offline operation by enabling devices to continue processing data even when they are disconnected from the internet.

upvoted 19 times

Appon Highly Voted 1 year, 10 months ago

Selected Answer: B

https://aws.amazon.com/blogs/machine-learning/anomaly-detection-with-amazon-sagemaker-edge-manager-using-aws-iot-greengrass-v2/

upvoted 5 times

career360guru Most Recent 1 year ago

Selected Answer: B

Option B

upvoted 1 times

dkcloudguru 1 year, 3 months ago

Option B: Greengrass supports offline operation

upvoted 1 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: B

Offline = IoT Greengrass

upvoted 2 times

SK_Tyagi 1 year, 4 months ago

If you can't commission your sensors

Consider the following questions.

Does the mobile phone running the Amazon Monitron App have a stable internet connection?

https://docs.aws.amazon.com/Monitron/latest/user-guide/troubleshooting.html

For commissioning a sensor, the mobile phone running the Amazon Monitron App should have internet connectivity.
upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

B for offline
upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: B

keyword = WS IoT Greengrass
upvoted 1 times

consultornetwork 1 year, 7 months ago

Selected Answer: B

Can't be D.
Amazon Monitron requires Internet connection.Q: Can I use Amazon Monitron when it is not connected to the AWS Region or in a disconnected environment?

A: Amazon Monitron Sensors and Gateways, and their use with the Amazon Monitron service, rely on connectivity over internet to the AWS Region.
https://aws.amazon.com/monitron/faqs/
Amazon Monitron Sensors and Gateways are not designed for disconnected operations or environments with no connectivity. We recommend that customers have highly available internet connectivity.
upvoted 3 times

Diego1414 1 year, 7 months ago

Selected Answer: B

AWS IoT Greengrass is software that extends cloud capabilities to local devices. This enables devices to collect and analyze data closer to the source of information, react autonomously to local events, and communicate securely with each other on local networks. Local devices can also communicate securely with AWS IoT Core and export IoT data to the AWS Cloud. AWS IoT Greengrass developers can use AWS Lambda functions and prebuilt connectors to create serverless applications that are deployed to devices for local execution.
upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: B

The ML model is run locally, so it can still provide feedback when the internet is down.
upvoted 3 times

hobokabobo 1 year, 9 months ago

Selected Answer: D

Quote "The company must be able to provide this feedback even if the factory's internet connectivity is down"
So everything that needs internet can be ignored. Leaves D.
While there is a lot of garbage text about how they process date with SargeMaker, the question only asks for a solution to detect failures in the equipment. Amazon Monitron does this plus it can work even when internet is down.

All other options provide solutions for things, the question didn't ask for and/or already in place and need internet.
upvoted 1 times

Untamables 1 year, 10 months ago

Selected Answer: B

The point is how to offload ML workloads to the local.
upvoted 2 times

Musk 1 year, 10 months ago

Selected Answer: B

Monitron is something different
upvoted 1 times

bititan 1 year, 11 months ago

Selected Answer: B

this is taking about detecting defects from an image that is taken from a camera. I would go for running a ML model on IoT greengras pc and transfer it to IoT core, then store it in s3 bucket, which can be called by api function via lambda to send it to users.
option D would monitor only sensor data of machines.

upvoted 4 times

schalke04 1 year, 11 months ago

Selected Answer: D

Amazon Monitron is a machine-learning based end-to-end condition monitoring system that detects potential failures within equipment. You can use it to implement a predictive maintenance program and reduce lost productivity from unplanned machine downtime. Amazon Monitron includes purpose-built sensors to capture vibration and temperature data, as well as gateways to automatically transfer data to the AWS Cloud. It also comes with an application in two versions. The mobile application handles system setup, analytics, and notification when tracking equipment conditions. The web application provides all the same functions as the mobile app except setup. Reliability managers can quickly deploy Amazon Monitron to track the machine health of industrial equipment, such as such as bearings, motors, gearboxes, and pumps, without any development work or specialized training.

upvoted 2 times

schalke04 1 year, 11 months ago

B is correct.

AWS IoT Greengrass enables ML inference locally using models that are created, trained, and optimized in the cloud using Amazon SageMaker, AWS Deep Learning AMI, or AWS Deep Learning Containers, and deployed on the edge devices

upvoted 3 times

youngprinceton 1 year, 10 months ago

when do you take the exam man i would like to see if everything is still valid after you test

upvoted 1 times

schalke04 1 year, 10 months ago

B is wrong, D is correct.

upvoted 2 times

Question #158 *Topic 1*

A solutions architect must create a business case for migration of a company's on-premises data center to the AWS Cloud. The solutions architect will use a configuration management database (CMDB) export of all the company's servers to create the case.

Which solution will meet these requirements MOST cost-effectively?

A. Use AWS Well-Architected Tool to import the CMDB data to perform an analysis and generate recommendations.

B. Use Migration Evaluator to perform an analysis. Use the data import template to upload the data from the CMDB export.

C. Implement resource matching rules. Use the CMDB export and the AWS Price List Bulk API to query CMDB data against AWS services in bulk.

D. Use AWS Application Discovery Service to import the CMDB data to perform an analysis.

Suggested Answer: *D*

Community vote distribution

B (89%) 11%

ZZ5 Highly Voted 1 year, 10 months ago

B

https://aws.amazon.com/blogs/architecture/accelerating-your-migration-to-aws/

Build a business case with AWS Migration Evaluator

The foundation for a successful migration starts with a defined business objective (for example, growth or new offerings). In order to enable the business drivers, the established business case must then be aligned to a technical capability (increased security and elasticity). AWS Migration Evaluator (formerly known as TSO Logic) can help you meet these objectives.

To get started, you can choose to upload exports from third-party tools such as Configuration Management Database (CMDB) or install a collector agent to monitor. You will receive an assessment after data collection, which includes a projected cost estimate and savings of running your on-premises workloads in the AWS Cloud. This estimate will provide a summary of the projected costs to re-host on AWS based on usage patterns. It will show the breakdown of costs by infrastructure and software licenses. With this information, you can make the business case and plan next steps.

upvoted 18 times

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: B

The AWS Migration Evaluator works by analyzing data about your current on-premises environment, including servers, storage, networking, and applications. It then provides a report that outlines the recommended AWS services and configurations that best match your existing infrastructure and applications. This report includes a detailed cost analysis that estimates the total cost of running your applications in the AWS cloud.

upvoted 11 times

liquen14 Most Recent 9 months, 4 weeks ago

This is again another example of completely stupid, nonsensical and useless exposition to ambiguity. Which one is correct because yeah, B seems to be well supported by https://aws.amazon.com/blogs/architecture/accelerating-your-migration-to-aws/
but in the faqs for AWS Application Discovery Service https://aws.amazon.com/blogs/architecture/accelerating-your-migration-to-aws/ there is literally a question about Application Discovery service

Q: Can I ingest data into Application Discovery Service from my existing configuration management database (CMDB)?

"Yes, you can import information about your on-premises servers and applications into the Migration Hub so you can track the status of application migrations. To import your data, you can download and populate the import CSV template and then upload it using the Migration Hub import console or by invoking the Application Discovery Service APIs"

So which one is correct? And what real knowledge are we getting from this pile of shit?

upvoted 2 times

saggy4 10 months, 3 weeks ago

Selected Answer: B

A - It is a questionnaire tool used to assess your AWS architecture

C - We will need to create Complex Application using SDK

D- Application Discovery is free and does support CMDB import but it can only give you plan and not a business use case
B - Correct answer: Free and helps you create bussiness use case.
upvoted 2 times

career360guru 1 year ago
Selected Answer: B
Option B
upvoted 1 times

bjexamprep 1 year ago
Selected Answer: B
Yes B is correct. But can you imagine any real architect in the world would trust such a solution for migration? It's a joke.
upvoted 1 times

joleneinthebackyard 1 year, 1 month ago
Selected Answer: B
When you see business case for migration, you think of Migration Evaluator.
According to ChatGPT,
A: AWS Well-Architected Tool: no option to import CMDB data
C: only provide insight about current data, doesnt consider the nuances of migration task
D: Application Discovery Service is for discover, not for building business cases
upvoted 1 times

bustedd 1 year, 2 months ago
Migration evaluation
B
upvoted 1 times

duriselvan 1 year, 3 months ago
https://www.youtube.com/watch?v=2qautbhuJC8
upvoted 1 times

Jonalb 1 year, 5 months ago
Selected Answer: D
D

This tools for Analitycs data : https://aws.amazon.com/pt/migration-evaluator/
Migration data or vm : https://aws.amazon.com/pt/application-discovery/faqs/
upvoted 2 times

NikkyDicky 1 year, 5 months ago
Selected Answer: B
B - use case for ME
upvoted 1 times

SkyZeroZx 1 year, 6 months ago
Selected Answer: B
Question say : Migration
then Anwser is : Migration Evaluator and other responde in this comments
upvoted 1 times

mfsec 1 year, 9 months ago
Selected Answer: B
B is the best fit
upvoted 3 times

kiran15789 1 year, 10 months ago
Selected Answer: B
Migration Evaluator is a complimentary service to create data-driven assessments and business cases for AWS cloud planning and migration.
upvoted 2 times

saurabh1805 1 year, 10 months ago
Selected Answer: B
B is right answer
upvoted 2 times

CloudFloater 1 year, 10 months ago

Selected Answer: B

B

Free service, focus on cost of migration

upvoted 3 times

spd 1 year, 10 months ago

Selected Answer: B

B - Evaluator

upvoted 2 times

Question #159 *Topic 1*

A company has a website that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Auto Scaling group. The ALB is associated with an AWS WAF web ACL.

The website often encounters attacks in the application layer. The attacks produce sudden and significant increases in traffic on the application server. The access logs show that each attack originates from different IP addresses. A solutions architect needs to implement a solution to mitigate these attacks.

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

A. Create an Amazon CloudWatch alarm that monitors server access. Set a threshold based on access by IP address. Configure an alarm action that adds the IP address to the web ACL's deny list.

B. Deploy AWS Shield Advanced in addition to AWS WAF. Add the ALB as a protected resource.

C. Create an Amazon CloudWatch alarm that monitors user IP addresses. Set a threshold based on access by IP address. Configure the alarm to invoke an AWS Lambda function to add a deny rule in the application server's subnet route table for any IP addresses that activate the alarm.

D. Inspect access logs to find a pattern of IP addresses that launched the attacks. Use an Amazon Route 53 geolocation routing policy to deny traffic from the countries that host those IP addresses.

Suggested Answer: *C*

Community vote distribution

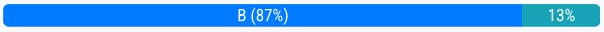

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: B

AWS Shield Advanced is focused on protecting against DDoS attacks, while AWS WAF is focused on protecting against web exploits. However, both services can be used together to provide comprehensive protection for your applications.

upvoted 14 times

shmoeee Most Recent 4 months, 3 weeks ago

Selected Answer: B

I chose B since this is a DDOS attack and also option A could cause issues if legitimate traffic gets thrown on the ACL deny list

upvoted 1 times

nelgeozcin 7 months, 2 weeks ago

Selected Answer: B

" The access logs show that each attack originates from different IP addresses. " implies DDoS

upvoted 1 times

Incognito013 10 months, 2 weeks ago

Selected Answer: A

Nothing mentioned about DDoS in the question, plus A is simplier and less operational oevrhead

upvoted 1 times

helloworldabc 10 months, 1 week ago

just B

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: B

Option B sounds most logical answer in terms of least operational overhead.
though it does not provide details about how to identify and add those IP addresses to Shield Advanced for DDos protection.

upvoted 2 times

Reejith 1 year, 7 months ago

I think its option A. Option B is a paid service and it is for DDoS. Here that attack is not DDoS and it is excess traffic generated at application layer by certain IPs. Not in a distributed attack pattern. Advanced shield will give DDoS+WAF. But you already have WAF and using which you can block the IPs

that is crossing set threshold. So option A is better choice. Option B is additional cost. Option C is wrong as you can not add deny rule in route table. Route table has only routes. Option D is operational overhead and then if you block the whole country , genuine traffic will also get blocked, which is not good.

upvoted 4 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: B

"Least" Operational Overhead - B

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

B 100%

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

Research more information and correct my answer

Letter B with this information

https://docs.aws.amazon.com/waf/latest/developerguide/ddos-app-layer-protections.html

upvoted 1 times

SkyZeroZx 2 years, 1 month ago

Selected Answer: A

For me it would be the letter A

Because AWS Shield Advanced is for DDOS attacks that happen at layer 3.

However, in the question they say attacks in the application layer

"The website often encounters attacks in the application layer."

For this reason, I would consider that it cannot be B and A would be a more feasible solution.

If anyone has more data, welcome to improve the community

Attached answer from Bard from Google

Here are some additional details about each solution:

upvoted 4 times

SkyZeroZx 2 years, 1 month ago

Solution C: This solution would require creating an AWS Lambda function, which is a paid service. AWS Lambda is a serverless compute service that allows you to run code without provisioning or managing servers. The Lambda function would be used to inspect access logs and identify IP addresses that are launching attacks. The function would then add those IP addresses to the application server's subnet route table, which would prevent traffic from those IP addresses from reaching the application server.

upvoted 1 times

SkyZeroZx 2 years, 1 month ago

Solution D: This solution would require inspecting access logs, which can be a time-consuming process. The access logs would be used to find a pattern of IP addresses that launched the attacks. The IP addresses could then be used to create a geolocation routing policy in Amazon Route 53. The geolocation routing policy would deny traffic from the countries that host those IP addresses.

Overall, solution A is the most efficient solution because it uses existing AWS services and does not require any additional infrastructure.

upvoted 1 times

SkyZeroZx 2 years, 1 month ago

Solution A: This solution is the most efficient because it uses existing AWS services and does not require any additional infrastructure. The CloudWatch alarm will monitor server access and trigger an action when the threshold is reached. The action can be configured to add the IP address to the web ACL's deny list, which will prevent traffic from that IP address from reaching the application server.

Solution B: This solution would require deploying AWS Shield Advanced, which is a paid service. AWS Shield Advanced provides additional protection against DDoS attacks, including application layer attacks. However, it is more expensive than AWS WAF.

upvoted 1 times

Daniel76 8 months, 3 weeks ago

The attack is at the application layer. Solution A detects attack by IP which is at network layer, hence it is not valid.

upvoted 1 times

dev112233xx 2 years, 2 months ago

Selected Answer: B

〃with the LEAST operational overhead〃 is AWS SHIELD Advanced without doubts✅

upvoted 3 times

hpipit 2 years, 3 months ago

Selected Answer: B

B 100% AWS SHIELD

upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: B

Deploy AWS Shield Advanced in addition to AWS WAF.

upvoted 2 times

rtgfdv3 2 years, 4 months ago

as long as i know or think to know, shield advanced, does nothing by default and needs to be configured.

https://docs.aws.amazon.com/waf/latest/developerguide/enable-ddos-prem.html
https://docs.aws.amazon.com/waf/latest/developerguide/getting-started-ddos.html
Note
Shield Advanced doesn't automatically protect your resources after you subscribe. You must specify the resources you want Shield Advanced to protect configure the protections.

upvoted 2 times

moota 2 years, 4 months ago

Selected Answer: B

According to ChatGPT, the ff are what you get with Advanced over Basic.

AWS Shield Advanced is a paid version of the service that provides additional protection against large scale and sophisticated DDoS attacks. This version includes all the features of the Basic version, but with additional capabilities such as 24/7 availability, a dedicated DDoS response team, and advanced attack analytics and reporting. Additionally, AWS Shield Advanced provides access to advanced DDoS protection and mitigation capabilities, such as the ability to customize protections for specific application requirements, and to mitigate attacks more quickly and effectively.

upvoted 3 times

Musk 2 years, 4 months ago

Selected Answer: B

Reading more about option B, I pick B

upvoted 4 times

Musk 2 years, 4 months ago

Not sure. With WAF you get Shield, which hs DDoS. Not sure the the Shield dvnced gives you much more.

upvoted 1 times

Question #160 *Topic 1*

A company has a critical application in which the data tier is deployed in a single AWS Region. The data tier uses an Amazon DynamoDB table and an Amazon Aurora MySQL DB cluster. The current Aurora MySQL engine version supports a global database. The application tier is already deployed in two Regions.

Company policy states that critical applications must have application tier components and data tier components deployed across two Regions. The RTO and RPO must be no more than a few minutes each. A solutions architect must recommend a solution to make the data tier compliant with company policy.

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

A. Add another Region to the Aurora MySQL DB cluster

B. Add another Region to each table in the Aurora MySQL DB cluster

C. Set up scheduled cross-Region backups for the DynamoDB table and the Aurora MySQL DB cluster

D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration

E. Use Amazon Route 53 Application Recovery Controller to automate database backup and recovery to the secondary Region

Suggested Answer: *BD*

Community vote distribution

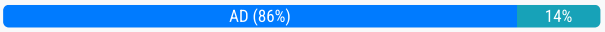

testingaws123 Highly Voted 1 year, 9 months ago

Badly written question:
"The RTO and RPO must be no more than a few minutes each."
What is few minutes mean? May be it is 2-3 min for me, may be it is 9-10 min for you.
upvoted 10 times

taer Highly Voted 1 year, 9 months ago

Selected Answer: AD
A. Add another Region to the Aurora MySQL DB cluster
D. Convert the existing DynamoDB table to a global table by adding another Region to its configuration
upvoted 5 times

career360guru Most Recent 1 year ago

Selected Answer: AD
A and D
upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: AD
A and D
upvoted 1 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: AD
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html
upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: AD
its AD
upvoted 1 times

pupsik 1 year, 6 months ago

Selected Answer: AD
For DynamoDB use global table, for Aurora use cross-region read-replicas.
upvoted 3 times

easytoo 1 year, 6 months ago

a-d-a-d-a-d-a-d-a-d

upvoted 1 times

Roontha 1 year, 6 months ago

Answer : A, D

https://docs.aws.amazon.com/prescriptive-guidance/latest/strategy-database-disaster-recovery/choosing-database.html

upvoted 1 times

God_Is_Love 1 year, 9 months ago

Selected Answer: AC

A solves multi region for DB layer. but question also asks for minimum RPO and RTO which means quick uptime of application in case of failure which is possible with backups.

https://aws.amazon.com/blogs/database/cost-effective-disaster-recovery-for-amazon-aurora-databases-using-aws-backup/

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide

/CrossRegionAccountCopyAWS.html

upvoted 3 times

SK_Tyagi 1 year, 4 months ago

Why use C and do replication with multiple steps when Global Tables support it

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html

upvoted 1 times

God_Is_Love 1 year, 9 months ago

Hint given is - Aurora MySQL engine version supports a global database which makes this possible -

https://d2908q01vomqb2.cloudfront.net/887309d048beef83ad3eabf2a79a64a389ab1c9f/2021/03/08/Aurora-Global-database-2.jpg

upvoted 4 times

schalke04 1 year, 11 months ago

Selected Answer: AD

A and D

upvoted 4 times

bititan 1 year, 11 months ago

Selected Answer: AD

you can create only db's not global tables, hence A and D

upvoted 4 times

Question #161 *Topic 1*

A telecommunications company is running an application on AWS. The company has set up an AWS Direct Connect connection between the company's on-premises data center and AWS. The company deployed the application on Amazon EC2 instances in multiple Availability Zones behind an internal Application Load Balancer (ALB). The company's clients connect from the on-premises network by using HTTPS. The TLS terminates in the ALB. The company has multiple target groups and uses path-based routing to forward requests based on the URL path.

The company is planning to deploy an on-premises firewall appliance with an allow list that is based on IP address. A solutions architect must develop a solution to allow traffic flow to AWS from the on-premises network so that the clients can continue to access the application.

Which solution will meet these requirements?

A. Configure the existing ALB to use static IP addresses. Assign IP addresses in multiple Availability Zones to the ALB. Add the ALB IP addresses to the firewall appliance.

B. Create a Network Load Balancer (NLB). Associate the NLB with one static IP addresses in multiple Availability Zones. Create an ALB-type target group for the NLB and add the existing ALAdd the NLB IP addresses to the firewall appliance. Update the clients to connect to the NLB.

C. Create a Network Load Balancer (NLB). Associate the LNB with one static IP addresses in multiple Availability Zones. Add the existing target groups to the NLB. Update the clients to connect to the NLB. Delete the ALB Add the NLB IP addresses to the firewall appliance.

D. Create a Gateway Load Balancer (GWLB). Assign static IP addresses to the GWLB in multiple Availability Zones. Create an ALB-type target group for the GWLB and add the existing ALB. Add the GWLB IP addresses to the firewall appliance. Update the clients to connect to the GWLB.

Suggested Answer: *A*

Community vote distribution

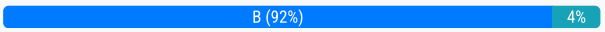

Untamables Highly Voted 1 year, 10 months ago

Selected Answer: B

The background is the below.
- The company is using ALB features and must keep them.
- The new on-premise firewall needs a static IP address of the ALB as the next hop.
- However, ALB cannot have a static IP address.
So the point is how ALB can have a static IP address endpoint.

Solution
https://aws.amazon.com/premiumsupport/knowledge-center/alb-static-ip/
upvoted 22 times

jojom19980 Highly Voted 1 year, 11 months ago

Selected Answer: B

it uses path-based routing to forward requests based on the URL path
upvoted 6 times

saggy4 Most Recent 10 months, 3 weeks ago

Selected Answer: B

A - Cannot assign static IP to ALB
C - Cannot attach target group directly as path-based forwarding is not possible with NLB
D - Gateway load balancer supports only Instance and IP as target
B - This is correct since using NLB we can have a static IP assigned and also attach ALB as target to NLB
upvoted 5 times

Spnohal 11 months, 1 week ago

https://aws.amazon.com/solutions/implementations/git-to-s3-using-webhooks/
upvoted 1 times

career360guru 1 year ago

Selected Answer: B

Option B is only feasible option is ALB is using path based routingg.
upvoted 1 times

CProgrammer 1 year ago
bjexamprep "Anyone help why A not correct?"
Where is the On Prem element, the Direct Connect, the ALB covering Multi AZ ?
"The objective of this question is achieved"
You don't even have the basic structure implemented
to attempt to address the questions requirements in your scenario
Regarding answer A :
https://repost.aws/knowledge-center/alb-static-ip
You can't assign a static IP address to an Application Load Balancer.
upvoted 1 times

bjexamprep 1 year ago
Selected Answer: A
Anyone can help explain why A is not correct? I created a private network facing ALB and it has a private IP address automatically created. Which means by adding the private IP address to the firewall, the objective of this question is achieved.
upvoted 2 times

saggy4 10 months, 3 weeks ago
A is not correct because, though the IP attached to the ALB is the private IP, the control of which IP is assign in with AWS, any change in the ALB can result in change of IP or even over a period of time AWS can change the IP (though it will be something in the CIDR)
upvoted 1 times

career360guru 1 year, 1 month ago
Selected Answer: B
Option B as ALB can not have static IP address so Option A is not possible.
upvoted 2 times

task_7 1 year, 3 months ago
D is also not the write answer
Target type
When you create a target group, you specify its target type, which determines how you specify its targets. After you create a target group, you cannot change its target type.

The following are the possible target types:

instance
The targets are specified by instance ID.

ip
The targets are specified by IP address.

When the target type is ip, you can specify IP addresses from one of the following CIDR blocks:

The subnets of the VPC for the target group

10.0.0.0/8 (RFC 1918)

100.64.0.0/10 (RFC 6598)

172.16.0.0/12 (RFC 1918)

192.168.0.0/16 (RFC 1918)
upvoted 1 times

task_7 1 year, 3 months ago
Elastic IP support
Network Load Balancer also allows you the option to assign an Elastic IP per Availability Zone (subnet) thereby providing your own fixed IP.
Both B anc C state single IP for multiple zones
upvoted 1 times

Gabehcoud 1 year, 4 months ago

Option B says "ALAdd" what is AL add? I see this very often. Can someone help to explain?
Create an ALB-type target group for the NLB and add the existing ALAdd the NLB IP addresses to the firewall appliance. Update the clients to connect to the NLB.

upvoted 1 times

khksoma 1 year, 5 months ago

A Gateway Load Balancer endpoint is a VPC endpoint that provides private connectivity between virtual appliances in the service provider VPC, and application servers in the service consumer VPC. The Gateway Load Balancer is deployed in the same VPC as that of the virtual appliances. These appliances are registered as a target group of the Gateway Load Balancer.
Since the firewall is deployed on-prem I dont think D is a viable option

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

B
need to keep ALB behind NLB for path routing

upvoted 1 times

Maria2023 1 year, 6 months ago

Selected Answer: B

Since ALB does not support static IP addresses by design then we need to use NLB before the ALB or instead. However, since we are heavily utilizing the application layer of the OSI then we cannot use NLB directly. Hence B remains the only choice

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: B

ALB's cannot use static IP's. NLB's have static IP's , addicionally need based on the URL path use ALB then
B is more apropiate

upvoted 1 times

rbm2023 1 year, 7 months ago

Selected Answer: B

I agree with B. since clients need access to the ALB using a private connection between on premises and AWS. The firewall which is inside company data center operates at network level but we cannot lose ALB due to many path based routing. So we need something like this:
https://www.scalefactory.com/blog/2021/12/13/aws-network-load-balancers-new-features/
https://www.scalefactory.com/blog/2021/12/13/aws-network-load-balancers-new-features/img/now-firewall-egress.png
and this:
https://aws.amazon.com/blogs/networking-and-content-delivery/application-load-balancer-type-target-group-for-network-load-balancer/

upvoted 3 times

God_Is_Love 1 year, 9 months ago

Selected Answer: D

https://aws.amazon.com/elasticloadbalancing/gateway-load-balancer/
Gateway Load Balancer helps you easily deploy, scale, and manage your third-party virtual appliances. It gives you one gateway for distributing traffic across multiple virtual appliances while scaling them up or down, based on demand. This decreases potential points of failure in your network and increases availability.

upvoted 1 times

God_Is_Love 1 year, 9 months ago

https://youtu.be/-j2smz_VCH4?t=1270
ALB (L7)- HTTP, HTTPS
NLB (L4)- TCP, UDP, TLS traffic
GWLB(L3)- IP traffic and 3rd party Appliances

upvoted 3 times

God_Is_Love 1 year, 9 months ago

AWS Gateway Load Balancer (GWLB) can terminate TLS traffic. GWLB supports SSL/TLS offloading, which means that it can terminate SSL/TLS connections from clients and then forward the decrypted traffic to backend servers over HTTP or HTTPS.

upvoted 1 times

Mickey321 1 year, 9 months ago

I think main question is can it support static IP address which is needed by the firmware to waitlist it?

upvoted 2 times

Question #162 *Topic 1*

A company runs an application on a fleet of Amazon EC2 instances that are in private subnets behind an internet-facing Application Load Balancer (ALB). The ALB is the origin for an Amazon CloudFront distribution. An AWS WAF web ACL that contains various AWS managed rules is associated with the CloudFront distribution.

The company needs a solution that will prevent internet traffic from directly accessing the ALB.

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

A. Create a new web ACL that contains the same rules that the existing web ACL contains. Associate the new web ACL with the ALB.

B. Associate the existing web ACL with the ALB.

C. Add a security group rule to the ALB to allow traffic from the AWS managed prefix list for CloudFront only.

D. Add a security group rule to the ALB to allow only the various CloudFront IP address ranges.

Suggested Answer: *D*

Community vote distribution

C (100%)

masssa Highly Voted 1 year, 10 months ago

Selected Answer: C

https://docs.amazonaws.cn/en_us/AmazonCloudFront/latest/DeveloperGuide/LocationsOfEdgeServers.html

AWS managed prefix list is more recommended.

upvoted 12 times

rbm2023 Highly Voted 1 year, 7 months ago

Selected Answer: C

https://docs.amazonaws.cn/en_us/AmazonCloudFront/latest/DeveloperGuide/LocationsOfEdgeServers.html

If your origin is hosted on Amazon and protected by an Amazon VPC security group, you can use the CloudFront managed prefix list to allow inbound traffic to your origin only from CloudFront's origin-facing servers, preventing any non-CloudFront traffic from reaching your origin

, imagine that your origin is an Amazon EC2 instance in the Europe (London) Region (eu-west-2). If the instance is in a VPC, you can create a security group rule that allows inbound HTTPS access from the CloudFront managed prefix list. This allows all of CloudFront's global origin-facing servers to reach the instance. If you remove all other inbound rules from the security group, you prevent any non-CloudFront traffic from reaching the instance

upvoted 5 times

career360guru Most Recent 1 year ago

Selected Answer: C

Option C

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: C

Option C

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

C for sure

upvoted 1 times

mfsec 1 year, 9 months ago

C. Add a security group rule to the ALB to allow traffic from the AWS managed prefix list for CloudFront only.

upvoted 2 times

ExamTopix01 1 year, 11 months ago

C https://aws.amazon.com/blogs/news/limit-access-to-your-origins-using-the-aws-managed-prefix-list-for-amazon-cloudfront/

upvoted 2 times

jojom19980 1 year, 11 months ago

Selected Answer: C

https://aws.amazon.com/about-aws/whats-new/2022/02/amazon-cloudfront-managed-prefix-list/

upvoted 3 times

Question #163 *Topic 1*

A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer. A recent security audit revealed that the company has configured encryption at rest for ElastiCache. However, the company did not configure ElastiCache to use encryption in transit. Additionally, users can access the cache without authentication.

A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption.

Which solution will meet these requirements?

A. Create an AUTH token. Store the token in AWS System Manager Parameter Store, as an encrypted parameter. Create a new cluster with AUTH, and configure encryption in transit. Update the application to retrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication.

B. Create an AUTH token. Store the token in AWS Secrets Manager. Configure the existing cluster to use the AUTH token, and configure encryption in transit. Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.

C. Create an SSL certificate. Store the certificate in AWS Secrets Manager. Create a new cluster, and configure encryption in transit. Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.

D. Create an SSL certificate. Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter. Update the existing cluster to configure encryption in transit. Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication.

Suggested Answer: *C*

Community vote distribution

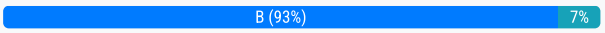

jimee11 1 month, 2 weeks ago

Selected Answer: B

Elasticache can be updated to support AUTH. Note: RBAC replaces AUTH now.

upvoted 1 times

zhen234 4 months, 3 weeks ago

Selected Answer: A

Encryption in transit cannot be enabled on an existing ElastiCache cluster. A new cluster must be created.

upvoted 2 times

d401c0d 4 months, 4 weeks ago

Selected Answer: B

Amazon ElastiCache for Redis now supports updates to encryption in transit on existing cluster resources. You can change the TLS configuration of your Redis clusters without re-building or re-provisioning them or impacting application availability. When enabling encryption in transit, your overall solution can remain connected to Redis clusters.

To get started, upgrade your Redis cluster to version 7 or above. You can then modify the encryption-in-transit property for your cluster using the Elasticache Console, API or CLI. This feature is available in all regions at no additional cost. To learn more, see the ElastiCache user guide.

upvoted 2 times

kylix75 5 months, 1 week ago

Selected Answer: A

The correct answer is A - Create an AUTH token, store it in Parameter Store, and create a new cluster with AUTH and in-transit encryption.
Key reasons:

ElastiCache doesn't allow enabling AUTH on existing clusters
SSL certificates aren't used for Redis authentication
Parameter Store is more cost-effective than Secrets Manager for this case
Solution meets both requirements: AUTH authentication and end-to-end encryption

upvoted 1 times

TewatiaAmit 8 months, 1 week ago

Selected Answer: A

A or B? Option B is suggesting to update the cluster which is not feasible. Once a cluster is created without encryption in transit, it cannot be modified to enable encryption in transit.

upvoted 1 times

Sin_Dan 8 months, 1 week ago

Selected Answer: A

Enabling encryption in transit on an existing ElastiCache cluster that wasn't originally configured with this feature is not possible. Encryption in transit, as well as encryption at rest, can only be specified at the time the cluster is created.

AWS Documentation on Encryption in Transit:

According to AWS ElastiCache documentation, if you want to enable encryption in transit, you must set this option when creating the ElastiCache cluster. Once a cluster is created without encryption in transit, it cannot be modified to enable this feature later. The same applies to Redis AUTH.

Thus, if a Redis cluster was deployed without encryption in transit, the only way to enable it is to create a new ElastiCache cluster with this setting enabled. Then, the data would need to be migrated from the existing cluster to the new one.

upvoted 3 times

JoeTromundo 8 months, 3 weeks ago

Selected Answer: B

B=Better :-)

upvoted 1 times

ke1dy 1 year, 1 month ago

Selected Answer: A

It seems to configure in-transit encryption in both new cluster and existing cluster, but updating is supported on Redis version 7 and later. So I will choose option A.

https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/in-transit-encryption.html#in-transit-encryption-constraints

upvoted 1 times

attila9778 6 months, 3 weeks ago

https://docs.aws.amazon.com/AmazonElastiCache/latest/dg/in-transit-encryption.html#in-transit-encryption-constraints

"Modifying the in-transit encryption setting, for an existing cluster, is supported on replication groups running Valkey 7.2 and later, and Redis OSS version 7 and later." => modifying is possible => so B

upvoted 1 times

helloworldabc 10 months, 1 week ago

just B

upvoted 3 times

gofavad926 1 year, 3 months ago

Selected Answer: B

A or B? I didn't read any comparison between these 2 options... For sure we need an auth token. Both, using SSM Parameter Store or Secrets Manager will work. Both, create a new cluster or update the current one will work. I will choose B because this approach avoids the need to set up a new cluster, potentially reducing effort and costs associated with migration or duplication of resources...

upvoted 3 times

career360guru 1 year, 6 months ago

Selected Answer: B

Option B

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

B, per redis docs.

EC encr in transit is a config option

upvoted 2 times

easytoo 2 years ago

b-b-b-b-b-b-b

Creating an AUTH token provides a form of authentication for accessing the ElastiCache cluster.
Storing the AUTH token in AWS Secrets Manager ensures secure and centralized management of the token.
Configuring the existing ElastiCache cluster to use the AUTH token enables authentication for accessing the cache.
Enabling encryption in transit ensures that data is encrypted when it is transferred between the client and the ElastiCache cluster.
Updating the application to retrieve the AUTH token from Secrets Manager and use it for authentication ensures that only authorized users can access the cache.

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: B

Create an AUTH token. Store the token in AWS Secrets Manager.

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: B

Redis CLI has AUTH command as a feature to SET/ROTATE strategies
https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/auth.html

upvoted 4 times

Zek 2 years, 3 months ago

B seems right.
To enable authentication on an existing Redis server, call the ModifyReplicationGroup API operation. Call ModifyReplicationGroup with the --auth-token parameter as the new token and the --auth-token-update-strategy with the value ROTATE.

After the modification is complete, the cluster supports the AUTH token specified in the auth-token parameter in addition to supporting connecting without authentication. Enabling authentication is only supported on Redis servers with encryption in transit (TLS) enabled.

https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/auth.html

upvoted 3 times

spd 2 years, 4 months ago

Selected Answer: B

As per https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/in-transit-encryption.html

upvoted 2 times

Question #164 *Topic 1*

A company is running a compute workload by using Amazon EC2 Spot Instances that are in an Auto Scaling group. The launch template uses two placement groups and a single instance type.

Recently, a monitoring system reported Auto Scaling instance launch failures that correlated with longer wait times for system users. The company needs to improve the overall reliability of the workload.

Which solution will meet this requirement?

A. Replace the launch template with a launch configuration to use an Auto Scaling group that uses attribute-based instance type selection.

B. Create a new launch template version that uses attribute-based instance type selection. Configure the Auto Scaling group to use the new launch template version.

C. Update the launch template Auto Scaling group to increase the number of placement groups.

D. Update the launch template to use a larger instance type.

Suggested Answer: *C*

Community vote distribution

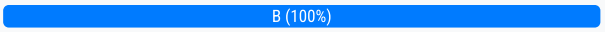

bititan Highly Voted 2 years, 4 months ago

Selected Answer: B

launch config is replaced by launch template hence is not advisible, option A rulled out. C is wrong because launch template cannot be updated. D is also wrong for the same reason

upvoted 14 times

Simon523 Highly Voted 1 year, 10 months ago

Selected Answer: B

As an alternative to manually specifying the instance types, you can specify the attributes that an instance must have, and Amazon EC2 will identify all the instance types with those attributes.

This is known as attribute-based instance type selection.

For example, you can specify the minimum and maximum number of vCPUs required for your instances, and EC2 Fleet will launch the instances using any available instance types that meet those vCPU requirements.

upvoted 6 times

JoeTromundo Most Recent 8 months, 3 weeks ago

Selected Answer: B

Correct answer: B

upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: B

Option B

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B

upvoted 1 times

totten 1 year, 8 months ago

Selected Answer: B

When you use attribute-based instance type selection, you allow AWS to diversify the instances across different instance types within a specified instance family or similar characteristics. This helps in reducing the risk of Spot Instance termination due to capacity issues or price fluctuations.

upvoted 5 times

rl97 1 year, 11 months ago

B

Amazon EC2 Auto Scaling can select from a wide range of instance types for launching Spot Instances. This meets the Spot best practice of being flexible about instance types, which gives the Amazon EC2 Spot service a better chance of finding and allocating your required amount of compute capacity.

upvoted 1 times

Christina666 1 year, 11 months ago

Selected Answer: B

key word "spot instance launch failure"-> attribute based selection

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: B

its a b

upvoted 1 times

easytoo 2 years ago

b-b-b-b-bb-b-

Creating a new launch template version allows for making changes to the template without disrupting the existing instances.
Using attribute-based instance type selection enables the Auto Scaling group to automatically select the most suitable instance type based on the defined attributes, such as availability zone, instance family, or instance size.
By leveraging attribute-based instance type selection, the Auto Scaling group can adapt to changing Spot Instance availability and launch instances in zones with higher availability, reducing launch failures.
Updating the launch template with this new version ensures that new instances launched by the Auto Scaling group utilize the improved instance selection process, thereby enhancing reliability.

upvoted 5 times

mfsec 2 years, 3 months ago

Selected Answer: B

B. Create a new launch template version that uses attribute-based instance type selection.

upvoted 2 times

Roontha 2 years, 1 month ago

Agreed with B

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: B

https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-instance-type-requirements.html#use-attribute-based-instance-type-selection-prerequisites

upvoted 2 times

kiran15789 2 years, 4 months ago

Selected Answer: B

Confused between B and D , will choose B

upvoted 1 times

saurabh1805 2 years, 4 months ago

Selected Answer: B

b is correct

https://aws.amazon.com/blogs/aws/new-attribute-based-instance-type-selection-for-ec2-auto-scaling-and-ec2-fleet/

upvoted 2 times

etechsystem_ts 2 years, 4 months ago

Selected Answer: B

B is correct

upvoted 1 times

Question #165 *Topic 1*

A company is migrating a document processing workload to AWS. The company has updated many applications to natively use the Amazon S3 API to store, retrieve, and modify documents that a processing server generates at a rate of approximately 5 documents every second. After the document processing is finished, customers can download the documents directly from Amazon S3.

During the migration, the company discovered that it could not immediately update the processing server that generates many documents to support the S3 API. The server runs on Linux and requires fast local access to the files that the server generates and modifies. When the server finishes processing, the files must be available to the public for download within 30 minutes.

Which solution will meet these requirements with the LEAST amount of effort?

A. Migrate the application to an AWS Lambda function. Use the AWS SDK for Java to generate, modify, and access the files that the company stores directly in Amazon S3.

B. Set up an Amazon S3 File Gateway and configure a file share that is linked to the document store. Mount the file share on an Amazon EC2 instance by using NFS. When changes occur in Amazon S3, initiate a RefreshCache API call to update the S3 File Gateway.

C. Configure Amazon FSx for Lustre with an import and export policy. Link the new file system to an S3 bucket. Install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS.

D. Configure AWS DataSync to connect to an Amazon EC2 instance. Configure a task to synchronize the generated files to and from Amazon S3.

Suggested Answer: *C*

Community vote distribution

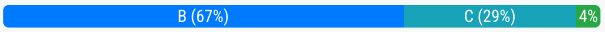

dev112233xx Highly Voted 2 years, 2 months ago

Selected Answer: B

B is correct imo

C is incorrect, FSx for Luster doesn't support NFS protocol

It actually support only POSIX protocol:

Custom (POSIX-compliant) protocol optimized for performance

upvoted 26 times

schalke04 Highly Voted 2 years, 4 months ago

Selected Answer: C

C:

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to access and process data concurrently from both a high-performance file system and from the S3 API.

upvoted 23 times

lxrdm 1 year, 12 months ago

I wouldnt choose Lustre.. would only pick it if its related to HPC (high performance computing), the amount of files generated here is nothing..

upvoted 5 times

rbm2023 2 years, 1 month ago

I disagree with option C. This is an example of how to mount a Lustre from an EC2 Linux system. it does not use NFS

sudo mount -t lustre <fsx-dns-name>@tcp:/<mount-point>

Amazon FSx for Lustre provides its own Lustre-specific mount command and protocol for mounting the file system on Linux instances.

The lustre file system type in the mount command indicates that it is specifically for mounting Lustre-based file systems, such as Amazon FSx for Lustre.

I would still go for option B

upvoted 9 times

d401c0d Most Recent 4 months, 4 weeks ago

Selected Answer: B

Luster does not support NFS.
upvoted 1 times

AWSum1 8 months, 3 weeks ago
Note that it keeps saying "The Server" implying 1 server and not a fleet or multiple. NFS is from 1 client to 1 server.

So C is incorrect
upvoted 1 times

JoeTromundo 8 months, 3 weeks ago
Selected Answer: B
Option C is not possible: how will you mount the document store on the EC2 instance through the Lustre client using NFS? Lustre is not compatible with NFS!
upvoted 1 times

xktm 10 months, 3 weeks ago
The English in this question is very confusing, what is it trying do? what is the problem? where is the processing server?
upvoted 6 times

duriselvan 1 year, 4 months ago
https://repost.aws/knowledge-center/storage-gateway-automate-refreshcache
upvoted 1 times

ninomfr64 1 year, 5 months ago
Selected Answer: B
A = migrating to lambda requires a lot of work and doesn't solve the need to have fast access to files
B = correct
C = FSx for Lustre doesn't support NFS
D = DataSynch can schedule transfer hourly, daily or weekly, cannot meet 30 minutes requirement
upvoted 7 times

career360guru 1 year, 6 months ago
Selected Answer: B
Option B as Fsx Luster though supports Linux, it does not support NFS.
upvoted 3 times

career360guru 1 year, 7 months ago
Selected Answer: B
B is right. Though it is meant to be used to with on-premise in Hybrid environment, it is possible to use it on EC2.
upvoted 3 times

Dougmaster 7 months, 4 weeks ago
B is right. If it wasn't possible to update that Linux server at that moment it implies they would have to remain it on premises for a while, in this case Amazon S3 File Gateway is the way to go.
upvoted 1 times

severlight 1 year, 7 months ago
Selected Answer: B
just because NFS mentioned with Lustre, but everything else is pointing to the Lustre: Linux, fast, read/writes to S3
upvoted 2 times

covabix879 1 year, 9 months ago
Selected Answer: C
B. Extra effort due to refreshCache API
D. DataSync runs in task schedule, which can't run faster than once per hour.
So remaining answer is C
upvoted 1 times

task_7 1 year, 9 months ago
Selected Answer: D
The core of the problem is make the file available in S3
for When the server finishes processing, the files must be available to the public for download within 30 minutes.
Which solution will meet these requirements with the LEAST amount of effort?
I think Option D (AWS DataSync) is a more straightforward and efficient choice.
upvoted 1 times

covabix879 1 year, 9 months ago

DataSync task cannot run faster than 1 hour. "Even with a cron expression, you can't schedule a task to run at an interval faster than 1 hour."
https://docs.aws.amazon.com/datasync/latest/userguide/task-scheduling.html

upvoted 5 times

Gabehcoud 1 year, 10 months ago

Selected Answer: B

The server is running Linux, How can we use Fsx?

upvoted 4 times

chikorita 1 year, 10 months ago

FSX for Lustre is for Linux and does not support Windows

upvoted 3 times

CloudHandsOn 1 year, 10 months ago

Selected Answer: B

I believe that B is correct, given that Lustre does not support NFS (it supports POSIX)

upvoted 3 times

xav1er 1 year, 10 months ago

Selected Answer: B

B as file gateway seems simple working solution for this. Lustre does not support NFS and might be an overkill for this solution - its primary used for HPC clusters. DataSync is rather for batch daad migrations and periodic data migration jobs, isn't it?

upvoted 5 times

softarts 1 year, 10 months ago

Selected Answer: B

don't understand the question and answer, include B&C. how does it mount to EC2 by using NFS? I think the processing server is running on Premise??

upvoted 3 times

Question #166 *Topic 1*

A delivery company is running a serverless solution in the AWS Cloud. The solution manages user data, delivery information, and past purchase details. The solution consists of several microservices. The central user service stores sensitive data in an Amazon DynamoDB table. Several of the other microservices store a copy of parts of the sensitive data in different storage services.

The company needs the ability to delete user information upon request. As soon as the central user service deletes a user, every other microservice must also delete its copy of the data immediately.

Which solution will meet these requirements?

A. Activate DynamoDB Streams on the DynamoDB table. Create an AWS Lambda trigger for the DynamoDB stream that will post events about user deletion in an Amazon Simple Queue Service (Amazon SQS) queue. Configure each microservice to poll the queue and delete the user from the DynamoDB table.

B. Set up DynamoDB event notifications on the DynamoDB table. Create an Amazon Simple Notification Service (Amazon SNS) topic as a target for the DynamoDB event notification. Configure each microservice to subscribe to the SNS topic and to delete the user from the DynamoDB table.

C. Configure the central user service to post an event on a custom Amazon EventBridge event bus when the company deletes a user. Create an EventBridge rule for each microservice to match the user deletion event pattern and invoke logic in the microservice to delete the user from the DynamoDB table.

D. Configure the central user service to post a message on an Amazon Simple Queue Service (Amazon SQS) queue when the company deletes a user. Configure each microservice to create an event filter on the SQS queue and to delete the user from the DynamoDB table.

Suggested Answer: *D*

Community vote distribution

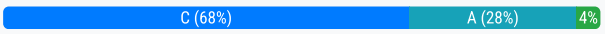

Untamables Highly Voted 2 years, 4 months ago

Selected Answer: A

The trigger is that the central user service deletes a user in the DynamoDB table. The DynamoDB Streams meets the requirement.
https://aws.amazon.com/blogs/database/how-to-perform-ordered-data-replication-between-applications-by-using-amazon-dynamodb-streams/
Option B is wrong. There is no feature named DynamoDB event notifications.

upvoted 17 times

Amac1979 2 years, 3 months ago

Correct, the point they want to make is central user service is system of record. You should not be deleting from other services until you delete from DynamoDB.

upvoted 1 times

kjcncjek 1 year, 10 months ago

how can you use 1 sqs queue for all microservises?

upvoted 3 times

jainparag1 1 year, 7 months ago

You can have many consumers which means any of the consumers can receive and process the message.

upvoted 5 times

CloudFloater Highly Voted 2 years, 4 months ago

Selected Answer: C

C seems correct; SQS is one queue to one microservice, could not find anything on dynamodb event notifications.

upvoted 17 times

jimee11 Most Recent 1 month, 2 weeks ago

Selected Answer: C

Poorly worded question. DynamoDB Streams is designed to do exactly what is required here. But, attempting multiple microservices reading the same SQS queue and updating the same table is wrong.

C is the best option of the mess.

upvoted 1 times

juanife 4 months, 2 weeks ago

it is MUCH MORE faster to send the event through eventbridge to microservices once the event of deletion needs to happen. I first read option A and thought it was the right one but have the microservices polling SQS QUEUE is less performant than the other one.
AND it's impossible for SQS to have multiple consumers as this is not the main purpose of this service, this is not a fan-out architecture with SQS and SNS.
Totally sure that C is the correct answer, I repeat, I thought it was A but it's not.

upvoted 1 times

chris_spencer 8 months, 3 weeks ago

Selected Answer: C

C, The problem with A the SQS solution ist that the "other microservices which stores data chunks seperatly". We do not know how many services are storing the userdata, and with SQS we would have one message on the queue which is processed by one of these microservices. how could the other microservices know that they have to delete the data when the message is allready consumed and processed?

upvoted 1 times

ry1999 9 months, 3 weeks ago

Selected Answer: C

SQS does not have a fan-out capability. You need SNS --> SQS to achieve the microservices to be notified. Hence A is incorrect and C is correct.

upvoted 3 times

Dgix 1 year, 3 months ago

Selected Answer: C

A is not viable since SQS is not used in a fan-out situation.
B is not viable since there's no such thing as "DynamoDB event notifications".
C is viable.
D is not viable, again due to the fact that SQS is not used for fan-out.

upvoted 4 times

career360guru 1 year, 6 months ago

Selected Answer: C

This is tricky question. C seems to be best and feasible. Rest options are not correct as they are using SQS where messages can be delivered only to one reader while in this scenario there are multiple microservices that needs to read the same message and delete the user information.

upvoted 4 times

CProgrammer 1 year, 6 months ago

Lets Ignore the insanity of
Several other microservices store in ---- different storage services. ------
central user service deletes a user, every other microservice must
also delete its copy of the data immediately.
YET ALL the options attempt a delete in the OG DynamoDB
Yeah OK Whatever Blue is green and Red is Orange these days.
BTW ans. == C , A will work but why poll SQS when Evt Brdg can invoke Microservice.
Personally I'd invoke a lambda to delete related records from the disparate data sources per KeyId and not bother the services but I'm not Architecting this mess maybe they want a clean log trail of the delete process as invoked by central user service whatever

upvoted 4 times

dankositzke 1 year, 4 months ago

Agreed. If this is an actual exam question, I am concerned about the intellect of the exam writers.

upvoted 4 times

Bad_Mat 1 year, 6 months ago

I vote for C because the question says: Delete the user IMMEDIATELY
A and D use SQS and messages in SQS can stay a pretty long time

upvoted 3 times

jainparag1 1 year, 7 months ago

Selected Answer: C

Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance, scalable storage for compute workloads. Powered by Lustre, the world's most popular high-performance file system, FSx for Lustre offers shared storage with sub-ms latencies, up to terabytes per second of throughput, and millions of IOPS. FSx for Lustre file systems can also be linked to Amazon Simple Storage Service (S3) buckets, allowing you to access and process data concurrently from both a high-performance file system and from the S3 API.

upvoted 1 times

jainparag1 1 year, 7 months ago

this is for Q165,

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C

upvoted 1 times

vjp_training 1 year, 9 months ago

Selected Answer: A

https://aws.amazon.com/vi/getting-started/hands-on/send-fanout-event-notifications/?nc1=f_ls

upvoted 2 times

Ganshank 1 year, 10 months ago

A real-world use case utterly destroyed with some of the worst possible options for solutions.

Simplest solution is to have the interested parties consume events off the DynamoDB streams and delete the user information in their respective datastores. Too many red herrings in the options given, and the only relatively sane one of the lot is Option C.

The bar for coming up with questions with SA professional keeps getting lowered.

upvoted 4 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: A

Event trigger from DynamoDb -- Choose DynamoDb Streams

upvoted 2 times

xav1er 1 year, 10 months ago

Where the hell is fan-out pattern? stupid answers ...

upvoted 3 times

aviathor 1 year, 10 months ago

* The central user service stores sensitive data in an Amazon DynamoDB table.

* Several of the other microservices store a copy of parts of the sensitive data in different storage services.

Apparently only the central user service stores user data in DynamoDB. The others use "different storage services". Yet, all of the answers focus on DynamoDB...

upvoted 1 times

Question #167 *Topic 1*

A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.

An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.

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

A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.

B. Allocate an Elastic IP address. Assign the Elastic IP address to the ALProvide the Elastic IP address to the customer.

C. Create an AWS Global Accelerator standard accelerator. Specify the ALB as the accelerator's endpoint. Provide the accelerator's IP addresses to the customer.

D. Configure an Amazon CloudFront distribution. Set the ALB as the origin. Ping the distribution's DNS name to determine the distribution's public IP address. Provide the IP address to the customer.

Suggested Answer: *B*

Community vote distribution

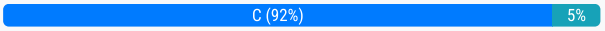

Untamables Highly Voted 2 years, 4 months ago

Selected Answer: C

https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html
Option A is wrong. AWS WAF does not support associating with NLB.
https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html
Option B is wrong. An ALB does not support an Elastic IP address.
https://aws.amazon.com/elasticloadbalancing/features/
upvoted 20 times

masssa Highly Voted 2 years, 4 months ago

static IP can made below method.
NLB (replace NLB from ALB)
NLB + ALB
global accelarator + ALB
original load balancer (ex. made by EC2 + nginx)
upvoted 18 times

AWSum1 Most Recent 8 months, 3 weeks ago

Selected Answer: C

Global Accelerator provides two global static public IPs that act as a fixed entry point to your application endpoints, such as Application Load Balancers, Network Load Balancers, Amazon Elastic Compute Cloud (EC2) instances, and elastic IPs.

https://aws.amazon.com/global-accelerator/
upvoted 1 times

ninomfr64 1 year, 5 months ago

Selected Answer: C

A = NLB doe not integrates with WAF
B = ALB cannot have Elastic IP attached, ALB cannot have static IP at all
C = corect
D = CloudFront distributions replies from many IPs, AWS manages a prefix list for this. Not easy to configure on customers side
upvoted 5 times

chsiri 6 months, 3 weeks ago

Why can't we create prefixlist with static ipaddress and assign to Cloudfront
upvoted 1 times

career360guru 1 year, 6 months ago

Selected Answer: C

Option C

upvoted 1 times

CProgrammer 1 year, 6 months ago

An Application Load Balancer cannot be assigned an Elastic IP address -- https://aws.amazon.com/blogs/networking-and-content-delivery/using-aws-lambda-to-enable-static-ip-addresses-for-application-load-balancers/

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C has least operational overhead. Option A is possible but changing ALB to NLB requires higher operational effort.

upvoted 2 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C - basic use case for GA

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

C. Create an AWS Global Accelerator standard accelerator.

upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

An Application Load Balancer cannot be assigned an Elastic IP address (static IP address).
https://stackoverflow.com/questions/55236806/how-to-assign-elastic-ip-to-application-load-balancer-in-aws

upvoted 1 times

God_Is_Love 2 years, 3 months ago

This feature allows you to migrate your applications to AWS without requiring your partners and customers to change their IP address whitelists. (which could be used in WAF)
BYOIP - Bring your own IP https://aws.amazon.com/blogs/networking-and-content-delivery/using-bring-your-own-ip-addresses-byoip-with-global-accelerator/

upvoted 2 times

kiran15789 2 years, 4 months ago

Selected Answer: C

https://aws.amazon.com/premiumsupport/knowledge-center/alb-static-ip/

Can assisng Static IP to ALB

upvoted 1 times

jojom19980 2 years, 4 months ago

Selected Answer: A

.....

upvoted 2 times

CloudInfrastructures 2 years, 4 months ago

C
WAF cannot be assoicated with NLB

upvoted 1 times

masssa 2 years, 4 months ago

NLB cannot be used when WAF is used

upvoted 1 times

ExamTopix01 2 years, 4 months ago

A
https://aws.amazon.com/jp/premiumsupport/knowledge-center/alb-static-ip/

upvoted 1 times

ExamTopix01 2 years, 4 months ago

Sorry C
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html
upvoted 1 times

schalke04 2 years, 4 months ago

This solution meets the requirement with the least operational overhead, as it only requires the allocation of an Elastic IP address, assignment to the ALB, and providing the address to the customer. The other options involve configuring additional services, which can increase operational overhead.
upvoted 1 times

bititan 2 years, 4 months ago

Selected Answer: C

this option has the least admin effort. A has more admin effort, B is not possible, D will not give static IP address
upvoted 4 times

schalke04 2 years, 4 months ago

Selected Answer: B

B will works
upvoted 1 times

Question #168 *Topic 1*

A company has a few AWS accounts for development and wants to move its production application to AWS. The company needs to enforce Amazon Elastic Block Store (Amazon EBS) encryption at rest current production accounts and future production accounts only. The company needs a solution that includes built-in blueprints and guardrails.

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

A. Use AWS CloudFormation StackSets to deploy AWS Config rules on production accounts.

B. Create a new AWS Control Tower landing zone in an existing developer account. Create OUs for accounts. Add production and development accounts to production and development OUs, respectively.

C. Create a new AWS Control Tower landing zone in the company's management account. Add production and development accounts to production and development OUs. respectively.

D. Invite existing accounts to join the organization in AWS Organizations. Create SCPs to ensure compliance.

E. Create a guardrail from the management account to detect EBS encryption.

F. Create a guardrail for the production OU to detect EBS encryption.

Suggested Answer: *BCE*

Community vote distribution

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: CDF

When you enable controls on an organizational unit (OU) that is registered with AWS Control Tower, preventive controls apply to all member accounts under the OU, enrolled and unenrolled. Detective controls apply to enrolled accounts only.
https://docs.aws.amazon.com/controltower/latest/userguide/controls.html

upvoted 13 times

Untamables Highly Voted 1 year, 10 months ago

Selected Answer: CDF

https://docs.aws.amazon.com/controltower/latest/userguide/controls.html
https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-encryption
AWS is now transitioning the previous term 'guardrail' new term 'control'.

upvoted 5 times

BelloMio Most Recent 2 months, 3 weeks ago

Selected Answer: CDE

I mean E is technically correct. The guardrail is created FROM the management account in Control Tower.

Even tho I would select F as well during the exam

upvoted 1 times

career360guru 1 year ago

Selected Answer: CDF

C, D, F are the right choices.

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: CDF

C, D, F

upvoted 1 times

bur4an 1 year, 3 months ago

Basically order is DCF of the setup

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: CDF

CDF for sure
upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: BCF

CEF
A) AWS Config not enforce rule
B) Why developer account ? is incorrect is management account
C) Sounds good
D) SCP for enforce sounds good
E) EBS encryption in managament account ? not only required in production
F) encryption in production OU sounds great
upvoted 3 times

SkyZeroZx 1 year, 6 months ago

CDF is correct
upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: BCF

https://www.examtopics.com/discussions/amazon/view/97939-exam-aws-certified-solutions-architect-professional-sap-c02/
upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: BCF

https://www.examtopics.com/discussions/amazon/view/97939-exam-aws-certified-solutions-architect-professional-sap-c02/
upvoted 1 times

Windows98 1 year, 6 months ago

Selected Answer: ACF

C because we want to use Control Tower

A and C because we're going to use Controls and Config

Not D because Control Tower is a parallel product to Organisations and it doesn't use SCPs although it can import existing OUs.
upvoted 3 times

Windows98 1 year, 6 months ago

I meant to say A and F because we're going to use Controls and Config
upvoted 1 times

Roontha 1 year, 7 months ago

Answer : C,D,F
upvoted 1 times

DWsk 1 year, 8 months ago

Selected Answer: ACF

I think the answer is ACF.
I don't think you need D once you have C. Also, control tower uses config rules to set up guardrails. See the link below:
https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#:~:text=isn%27t%20enabled%20on%20any%20OUs.-,The%20artifact%20for%20this%20control%20is%20the%20following%20AWS%20Config%2AWSTemplateFormatVersion%3A%202010%2D09%2D09
upvoted 2 times

xenodamus 1 year, 7 months ago

You still need to invite accounts before you can organize them in OUs. All steps are needed. I don't like the way they scatter between answers though.
upvoted 2 times

mfsec 1 year, 9 months ago

Selected Answer: CDF

CDF seems the best choice
upvoted 1 times

dummy1777 1 year, 10 months ago

B. Create a new AWS Control Tower landing zone in an existing developer account. Create OUs for accounts. Add production and development accounts to production and development OUs, respectively.
D. Invite existing accounts to join the organization in AWS Organizations. Create SCPs to ensure compliance.
F. Create a control for the production OU to detect EBS encryption.

By creating a new AWS Control Tower landing zone, the company can create OUs for accounts and add them to the appropriate production and development OUs. This will enable centralized governance and enforce consistent policies and best practices. The company can then invite existing accounts to join the organization in AWS Organizations and create SCPs to ensure compliance. Finally, the company can create a control for the production OU to detect EBS encryption, ensuring that encryption at rest is enforced in production accounts.
upvoted 2 times

spd 1 year, 10 months ago

Selected Answer: CDF

Answer is CDF
https://docs.aws.amazon.com/controltower/latest/userguide/controls.html
https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-encryption
upvoted 1 times

c73bf38 1 year, 10 months ago

The artifact for this control is AWS Config rule and AWS Config rules cannot be deployed using AWS CloudFormation StackSets.
upvoted 1 times

c73bf38 1 year, 10 months ago

moderator, delete above as the statement is incorrect that I posted, don't approve post.
upvoted 1 times

Musk 1 year, 10 months ago

Selected Answer: ABD

In F, guardrails are proposed to detect. Guardrails don't detect but prevent.
upvoted 1 times

Musk 1 year, 10 months ago

I found this, and after further reading I vote for CDF: https://docs.aws.amazon.com/controltower/latest/userguide/strongly-recommended-controls.html#ebs-enable-encryption
upvoted 1 times

Question #169 *Topic 1*

A company is running a critical stateful web application on two Linux Amazon EC2 instances behind an Application Load Balancer (ALB) with an Amazon RDS for MySQL database. The company hosts the DNS records for the application in Amazon Route 53. A solutions architect must recommend a solution to improve the resiliency of the application.

The solution must meet the following objectives:

• Application tier: RPO of 2 minutes. RTO of 30 minutes
• Database tier: RPO of 5 minutes. RTO of 30 minutes

The company does not want to make significant changes to the existing application architecture. The company must ensure optimal latency after a failover.

Which solution will meet these requirements?

A. Configure the EC2 instances to use AWS Elastic Disaster Recovery. Create a cross-Region read replica for the RDS DB instance. Create an ALB in a second AWS Region. Create an AWS Global Accelerator endpoint, and associate the endpoint with the ALBs. Update DNS records to point to the Global Accelerator endpoint.

B. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes. Configure RDS automated backups. Configure backup replication to a second AWS Region. Create an ALB in the second Region. Create an AWS Global Accelerator endpoint, and associate the endpoint with the ALBs. Update DNS records to point to the Global Accelerator endpoint.

C. Create a backup plan in AWS Backup for the EC2 instances and RDS DB instance. Configure backup replication to a second AWS Region. Create an ALB in the second Region. Configure an Amazon CloudFront distribution in front of the ALB. Update DNS records to point to CloudFront.

D. Configure the EC2 instances to use Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes. Create a cross-Region read replica for the RDS DB instance. Create an ALB in a second AWS Region. Create an AWS Global Accelerator endpoint, and associate the endpoint with the ALBs.

Suggested Answer: *B*

Community vote distribution

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: A

DRS includes EC2 instances as well not just data related as offered by DLM or Backup

Q: What operating systems and applications are supported by AWS DRS?
A: You can use AWS DRS to recover all of your applications and databases that run on supported Windows and Linux operating system versions. This includes critical databases such as Oracle, MySQL, and SQL Server, and enterprise applications such as SAP.

AWS Elastic Disaster Recovery (DRS) vs AWS DLM vs AWS Backup

You should use DLM when you want to automate the creation, retention, and deletion of EBS snapshots. You should use AWS Backup to manage and monitor backups across the AWS services you use, including EBS volumes, from a single place.

upvoted 23 times

bititan Highly Voted 2 years, 4 months ago

Selected Answer: A

its understood that others cannot meet the RTO and RPO requirements, because restore from back can take time based on the size of the data

upvoted 11 times

sarlos Most Recent 1 year, 1 month ago

Why not C?

upvoted 1 times

helloworldabc 10 months, 1 week ago

just A
upvoted 1 times

tushar321 1 year, 2 months ago
DRS Maintains state of EC2 machines while snapshot doesnt
upvoted 1 times

career360guru 1 year, 6 months ago
Selected Answer: A
Option A
upvoted 1 times

career360guru 1 year, 7 months ago
Selected Answer: A
Option A
upvoted 1 times

DiaaCloud 1 year, 8 months ago
A is correct
D is not correct because snapshot is one region and must to be copied and keep in sync to DR region which cannot meet the RTO...for sure D is wrong
upvoted 1 times

nharaz 1 year, 9 months ago
Selected Answer: A
DRS is faster to recover than Backups > https://youtu.be/07EHsPuKXc0?si=w_dZQKOAynE2T4JY
upvoted 1 times

NikkyDicky 1 year, 12 months ago
Selected Answer: A
A for low RPO
upvoted 1 times

Jesuisleon 2 years, 1 month ago
I don't understand the sentence "Update DNS records to point to the Global Accelerator endpoint" in A and B. It doesn't make sense. I think it should "update DNS records to point to the GA two static IP addresses or GA's DNS name
upvoted 1 times

dev112233xx 2 years, 2 months ago
Selected Answer: A
RDS Cross-region replication has the best RPO and RTO:
https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/

https://docs.aws.amazon.com/prescriptive-guidance/latest/strategy-database-disaster-recovery/choosing-database.html

AWS Elastic Disaster Recovery also provide the best RTO/RPO (with Warm standby and active-active)
https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/rel_planning_for_recovery_disaster_recovery.html
upvoted 5 times

OCHT 2 years, 2 months ago
Selected Answer: D
You are correct that AWS Elastic Disaster Recovery (DRS) can be used to recover both data and EC2 instances. However, in the scenario described in the question, the specified RPO and RTO objectives for the application tier can be met using Amazon Data Lifecycle Manager (Amazon DLM) to take snapshots of the EBS volumes attached to the EC2 instances.

While restoring from a backup can take time depending on the size of the data, using Amazon DLM to take snapshots of the EBS volumes provides a way to recover data within the specified RPO of 2 minutes and RTO of 30 minutes for the application tier.

In addition, creating a cross-Region read replica for the RDS DB instance provides a way to recover data within the specified RPO of 5 minutes and RTO of 30 minutes for the database tier.
upvoted 2 times

michele_scar 1 year, 1 month ago
Option D doesn't mention DNS, so it's not correct
upvoted 1 times

OCHT 2 years, 2 months ago

Overall, while AWS Elastic Disaster Recovery (DRS) can be a useful service in certain scenarios, it is not necessary in this case because the specified RPO and RTO objectives can be met using other AWS services such as Amazon Data Lifecycle Manager (Amazon DLM) and cross-Region read replicas for the RDS DB instance.

upvoted 1 times

BasselBuzz 1 year, 11 months ago

The process of starting up new instances and mount the EBS volumes to them will absolutely take more than 30 minutes.

upvoted 1 times

OCHT 2 years, 2 months ago

Overall, while AWS Elastic Disaster Recovery (DRS) can be a useful service in certain scenarios, it is not necessary in this case because the specified RPO and RTO objectives can be met using other AWS services such as Amazon Data Lifecycle Manager (Amazon DLM) and cross-Region read replicas for the RDS DB instance.

upvoted 1 times

OCHT 2 years, 2 months ago

Option A is not the best solution because it involves using AWS Elastic Disaster Recovery, which is not necessary to meet the specified RPO and RTO objectives for the application and database tiers.

AWS Elastic Disaster Recovery is a service that helps customers prepare for and recover from disasters by providing a cost-effective, fully managed, and scalable solution for disaster recovery. While it can be useful in certain scenarios, it is not necessary in this case because the specified RPO and RTO objectives can be met using other AWS services such as Amazon Data Lifecycle Manager (Amazon DLM) and cross-Region read replicas for the RDS DB instance.

Therefore, Option D is a better solution because it meets the specified requirements without introducing unnecessary complexity or cost.

upvoted 1 times

Musk 2 years, 4 months ago

Selected Answer: A

I agree it's A

upvoted 2 times

schalke04 2 years, 4 months ago

Selected Answer: A

DRS should fulfill the requirements

upvoted 3 times

Question #170 *Topic 1*

A solutions architect wants to cost-optimize and appropriately size Amazon EC2 instances in a single AWS account. The solutions architect wants to ensure that the instances are optimized based on CPU, memory, and network metrics.

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

A. Purchase AWS Business Support or AWS Enterprise Support for the account.

B. Turn on AWS Trusted Advisor and review any “Low Utilization Amazon EC2 Instances” recommendations.

C. Install the Amazon CloudWatch agent and configure memory metric collection on the EC2 instances.

D. Configure AWS Compute Optimizer in the AWS account to receive findings and optimization recommendations.

E. Create an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest.

Suggested Answer: *BD*

Community vote distribution

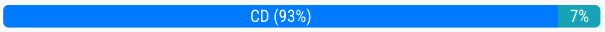

God_Is_Love Highly Voted 2 years, 3 months ago

Selected Answer: CD

Not B because, Trusted Advisor is available for Enterprise support only which is not cheap and the SA needs to cost optimize here. CPU, memory, and network relate to Compute so D for sure. C will enable to know how much actual memory/CPU is needed for instances and SA can provision based on cw logs

upvoted 11 times

LuongTo Most Recent 7 months ago

I would go for CD

B is more "CPU utilization

C is more "memory metrics"

D is more CPU and network metrics

then CD is more comprehensive while DB miss the "memory" part

upvoted 1 times

JoeTromundo 8 months, 3 weeks ago

Selected Answer: CD

Correct answer: C and D

"Memory utilization metrics are analyzed for the following resources: EC2 instances with the CloudWatch agent that's installed on them."

upvoted 1 times

TonytheTiger 1 year, 3 months ago

Selected Answer: CD

NOT Option B - To have Compute Optimizer analyze the memory utilization metric of your instances, install the CloudWatch agent on your instances. Enabling Compute Optimizer to analyze memory utilization data for your instances provides an additional measurement of data that further improves Compute Optimizer's recommendations.

https://docs.aws.amazon.com/compute-optimizer/latest/ug/metrics.html#ec2-metrics-analyzed

upvoted 2 times

career360guru 1 year, 6 months ago

Selected Answer: CD

Option C and D

upvoted 2 times

AWSStudyBuddy 1 year, 6 months ago

The solutions architect should take the following two steps to meet the requirements:

Configure AWS Compute Optimizer in the AWS account to receive findings and optimization recommendations. Compute Optimizer uses machine learning to analyze historical utilization metrics and provides recommendations to reduce costs and increase workload performance by

recommending the optimal instance types.

Turn on AWS Trusted Advisor and review any "Low Utilization Amazon EC2 Instances" recommendations. Trusted Advisor checks for underutilized instances and provides recommendations to right-size them, helping optimize costs.

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: CD

C and D

upvoted 2 times

Russs99 1 year, 9 months ago

Selected Answer: BD

AWS Trusted Advisor and AWS Compute Optimizer can both provide recommendations for right-sizing EC2 instances without requiring the installation of the CloudWatch agent or the collection of memory metrics.

The CloudWatch agent is primarily used for monitoring EC2 instances and collecting data for performance analysis. While it can be helpful to collect memory metrics for EC2 instances, it is not required for cost-optimizing and appropriately sizing them.

upvoted 3 times

Simon523 1 year, 10 months ago

Selected Answer: CD

AWS Compute Optimizer recommends optimal AWS resources for your workloads to reduce costs and improve performance by using machine learning to analyze historical utilization metrics.

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: CD

Cloud Watch Agent for memory metric & Compute Optimizer for recommendations

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: CD

cd for sure

upvoted 1 times

iamunstopable 2 years, 2 months ago

A & B will incur more cost. CD are correct

upvoted 2 times

Roontha 2 years, 1 month ago

Agreed. Answers are C,D

https://docs.aws.amazon.com/compute-optimizer/latest/ug/metrics.html

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: CD

CD is right

upvoted 1 times

saurabh1805 2 years, 4 months ago

Selected Answer: CD

trusted advisor does not take memory in consideration hence CD is right answer.

https://docs.aws.amazon.com/awssupport/latest/user/cost-optimization-checks.html

upvoted 1 times

CloudFloater 2 years, 4 months ago

D,OK.. but, why not B trusted advisor rather than C cloudwatch ?

upvoted 1 times

hobokabobo 2 years, 3 months ago

Memory taken by the os is almost always 100% - but most of it caches, buffers. To get you need the actually used memory by applications. This is number is os specific(need to ask the os how the memory is used: only caches or actual use?) and as such can't be gathered from the virtualizer. So you need an agent for that.

upvoted 1 times

rtgfdv3 2 years, 4 months ago

seems like you need cloud watch agent installed in order to check memory parameter

Note:

To have Compute Optimizer analyze the memory utilization of your instances, install the CloudWatch agent on your instances. Enabling Compute Optimizer to analyze memory utilization data for your instances provides an additional measurement of data that further improves Compute Optimizer's recommendations

https://docs.aws.amazon.com/compute-optimizer/latest/ug/metrics.html

upvoted 3 times

Musk 2 years, 4 months ago

Selected Answer: CD

CD according to https://docs.aws.amazon.com/compute-optimizer/latest/ug/metrics.html

upvoted 2 times

spd 2 years, 4 months ago

Selected Answer: CD

For Memory - CLoudwatch and

Compute Optimizer

upvoted 3 times

c73bf38 2 years, 4 months ago

What about the other metrics?

CPU and network metrics.

upvoted 1 times

c73bf38 2 years, 4 months ago

CD is correct, cloudwatch agents supports the metrics mentioned.

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/metrics-collected-by-CloudWatch-agent.html

upvoted 2 times

Question #171 *Topic 1*

A company uses an AWS CodeCommit repository. The company must store a backup copy of the data that is in the repository in a second AWS Region.

Which solution will meet these requirements?

A. Configure AWS Elastic Disaster Recovery to replicate the CodeCommit repository data to the second Region.

B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule. Create a cross-Region copy in the second Region.

C. Create an Amazon EventBridge rule to invoke AWS CodeBuild when the company pushes code to the repository. Use CodeBuild to clone the repository. Create a .zip file of the content. Copy the file to an S3 bucket in the second Region.

D. Create an AWS Step Functions workflow on an hourly schedule to take a snapshot of the CodeCommit repository. Configure the workflow to copy the snapshot to an S3 bucket in the second Region

Suggested Answer: *C*

Community vote distribution

bjexamprep Highly Voted 1 year, 6 months ago

Selected Answer: C

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automate-event-driven-backups-from-codecommit-to-amazon-s3-using-codebuild-and-cloudwatch-events.html

Hard to believe a product from AWS can be designed in such an amateur way.

upvoted 12 times

GabrielShiao 5 months, 2 weeks ago

It is unbelievable for such a solution. In particular, it happens in the company like AWS

upvoted 1 times

nimbus_00 Most Recent 7 months, 1 week ago

Selected Answer: C

Yeah...Deprecating CodeCommit was the right decision!

upvoted 1 times

AWSum1 8 months, 3 weeks ago

AWS Backup does not support AWS CodeCommit directly.

upvoted 2 times

AWSum1 8 months, 3 weeks ago

C is correct

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C

upvoted 2 times

severlight 1 year, 7 months ago

Selected Answer: C

yes, AWS Backup cannot do this for you, so you should use Code Build to clone repo and upload zip to s3

upvoted 3 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

its a C

upvoted 1 times

easytoo 2 years ago

b-b-b-b-b-b-b-b

upvoted 1 times

easytoo 2 years ago

b in incorrect as AWS Backup does not backup code commit as a source.

upvoted 3 times

easytoo 2 years ago

c-c-c-c-cc-c-c-c-c-c-c

upvoted 2 times

Roontha 2 years, 1 month ago

Answer : C

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automate-event-driven-backups-from-codecommit-to-amazon-s3-using-codebuild-and-cloudwatch-events.html

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: C

C for sure

upvoted 2 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automate-event-driven-backups-from-codecommit-to-amazon-s3-using-codebuild-and-cloudwatch-events.html

upvoted 1 times

kiran15789 2 years, 4 months ago

Selected Answer: C

https://www.automat-it.com/post/backup-aws-codecommit

upvoted 3 times

c73bf38 2 years, 4 months ago

Selected Answer: C

C is correct, AWS Backup does not backup code commit as a source.

upvoted 2 times

lunt 2 years, 4 months ago

Selected Answer: C

B is wrong > AWS Backup does not support CodeCommit as source.

A is out.

C is right.

upvoted 2 times

Musk 2 years, 4 months ago

Selected Answer: C

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automate-event-driven-backups-from-codecommit-to-amazon-s3-using-codebuild-and-cloudwatch-events.html

upvoted 2 times

c73bf38 2 years, 4 months ago

Selected Answer: B

It says backup so I think B is the answer:

B. Use AWS Backup to back up the CodeCommit repository on an hourly schedule. Create a cross-Region copy in the second Region.

upvoted 1 times

c73bf38 2 years, 4 months ago

Changing to C, thanks.

upvoted 2 times

spd 2 years, 4 months ago

Selected Answer: C

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/deploy-code-in-multiple-aws-regions-using-aws-codepipeline-aws-codecommit-and-aws-codebuild.html

https://medium.com/geekculture/replicate-aws-codecommit-repositories-between-regions-using-codebuild-and-codepipeline-39f6b8fcefd2

upvoted 4 times

Question #172 *Topic 1*

A company has multiple business units that each have separate accounts on AWS. Each business unit manages its own network with several VPCs that have CIDR ranges that overlap. The company's marketing team has created a new internal application and wants to make the application accessible to all the other business units. The solution must use private IP addresses only.

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

A. Instruct each business unit to add a unique secondary CIDR range to the business unit's VPC. Peer the VPCs and use a private NAT gateway in the secondary range to route traffic to the marketing team.

B. Create an Amazon EC2 instance to serve as a virtual appliance in the marketing account's VPC. Create an AWS Site-to-Site VPN connection between the marketing team and each business unit's VPC. Perform NAT where necessary.

C. Create an AWS PrivateLink endpoint service to share the marketing application. Grant permission to specific AWS accounts to connect to the service. Create interface VPC endpoints in other accounts to access the application by using private IP addresses.

D. Create a Network Load Balancer (NLB) in front of the marketing application in a private subnet. Create an API Gateway API. Use the Amazon API Gateway private integration to connect the API to the NLB. Activate IAM authorization for the API. Grant access to the accounts of the other business units.

Suggested Answer: *D*

Community vote distribution

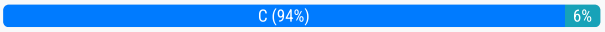

spd Highly Voted 2 years, 4 months ago

Selected Answer: C

Private link is the solution for IP Overlapping and Securely access the app between accounts

upvoted 16 times

c73bf38 Highly Voted 2 years, 4 months ago

Selected Answer: C

With AWS PrivateLink, the marketing team can create an endpoint service to share their internal application with other accounts securely using private IP addresses. They can grant permission to specific AWS accounts to connect to the service and create interface VPC endpoints in the other accounts to access the application by using private IP addresses. This option does not require any changes to the network of the other business units, and it does not require peering or NATing. This solution is both scalable and secure.

upvoted 11 times

alexsanteeno Most Recent 1 year, 6 months ago

Selected Answer: B

"LEAST OPERATIONAL OVERHEAD" - is key word in a question. Its not so easy to migrate any on-premise infra to any AWS. Looking at the answers here I see no one eve done that before and just answering as from AWS docs.

The easiest way to migrate any on-premise infra - ec2

upvoted 1 times

helloworldabc 10 months, 1 week ago

just C

upvoted 2 times

StevePace 1 year, 3 months ago

who mentioned migration?!

upvoted 1 times

honoga4853 1 year, 6 months ago

Selected Answer: B

"LEAST OPERATIONAL OVERHEAD" - is key word in a question. Its not so easy to migrate any on-premise infra to any AWS. Looking at the answers here I see no one eve done that before and just answering as from AWS docs.

The easiest way to migrate any on-premise infra - ec2

upvoted 1 times

helloworldabc 10 months, 1 week ago

just C
upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C
upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C for sure
upvoted 1 times

Alabi 2 years ago

Selected Answer: C

The solution that will meet the requirements with the least operational overhead is:

C. Create an AWS PrivateLink endpoint service to share the marketing application. Grant permission to specific AWS accounts to connect to the service. Create interface VPC endpoints in other accounts to access the application using private IP addresses.

AWS PrivateLink provides secure and scalable private connectivity between VPCs, AWS services, and on-premises applications, without using public IP addresses. In this case, you can create an AWS PrivateLink endpoint service for the marketing application, which allows other business units to access the application using private IP addresses.

By granting permission to specific AWS accounts to connect to the PrivateLink endpoint service, you can control access to the marketing application. Then, in each business unit's VPC, you can create interface VPC endpoints to connect to the PrivateLink service, allowing them to access the marketing application privately.
upvoted 2 times

mfsec 2 years, 3 months ago

Selected Answer: C

Private link
upvoted 1 times

God_Is_Love 2 years, 3 months ago

Selected Answer: C

Networking & Content Delivery blog -
https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/
upvoted 5 times

Question #173 *Topic 1*

A company needs to audit the security posture of a newly acquired AWS account. The company's data security team requires a notification only when an Amazon S3 bucket becomes publicly exposed. The company has already established an Amazon Simple Notification Service (Amazon SNS) topic that has the data security team's email address subscribed.

Which solution will meet these requirements?

A. Create an S3 event notification on all S3 buckets for the isPublic event. Select the SNS topic as the target for the event notifications.

B. Create an analyzer in AWS Identity and Access Management Access Analyzer. Create an Amazon EventBridge rule for the event type "Access Analyzer Finding" with a filter for "isPublic: true." Select the SNS topic as the EventBridge rule target.

C. Create an Amazon EventBridge rule for the event type "Bucket-Level API Call via CloudTrail" with a filter for "PutBucketPolicy." Select the SNS topic as the EventBridge rule target.

D. Activate AWS Config and add the cloudtrail-s3-dataevents-enabled rule. Create an Amazon EventBridge rule for the event type "Config Rules Re-evaluation Status" with a filter for "NON_COMPLIANT." Select the SNS topic as the EventBridge rule target.

Suggested Answer: *A*

Community vote distribution

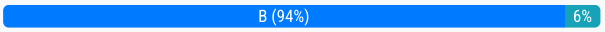

dkx Highly Voted 1 year, 5 months ago

A. No, because Amazon S3 can NOT currently publish notifications for isPublic events.
https://docs.aws.amazon.com/AmazonS3/latest/userguide/EventNotifications.html

B. Yes, because IAM Access Analyzer for S3 alerts you to S3 buckets that are configured to allow access to anyone on the internet or other AWS accounts
https://aws.amazon.com/blogs/security/how-to-prioritize-iam-access-analyzer-findings/

C. No, because PutBucketPolicy notifies us of an Amazon S3 bucket policy event to an Amazon S3 bucket, and we are looking for a SPECIFIC event to the bucket permissions, not ALL events.

D. No, because cloudtrail-s3-dataevents-enabled checks if at least one AWS CloudTrail trail is logging Amazon Simple Storage Service (Amazon S3) data events for all S3 buckets.

https://docs.aws.amazon.com/config/latest/developerguide/cloudtrail-s3-dataevents-enabled.html
upvoted 14 times

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: B
https://docs.aws.amazon.com/AmazonS3/latest/userguide/access-analyzer.html
upvoted 12 times

God_Is_Love 1 year, 9 months ago

Click on the "Create rule" button.

Enter a name for the rule and a brief description, if desired.

Under "Define pattern", select "Event pattern".

Select "Custom pattern".

In the "Event pattern" field, enter the following code:

{
"source": ["aws.securityhub"],
"detail-type": ["Access Analyzer Finding"],

"detail": {
"findings": [
{
"isPublic": [
true
]
}
]
}
}

This code will match all Access Analyzer Finding events where the "isPublic" field is set to "true".
upvoted 8 times

AimarLeo Most Recent 11 months ago
This question.. is seriously ! a googling one
upvoted 1 times

dkcloudguru 1 year, 3 months ago
Option B
upvoted 1 times

NikkyDicky 1 year, 5 months ago
Selected Answer: B
it's B
upvoted 2 times

Maria2023 1 year, 6 months ago
Selected Answer: B
Ideally, I would use config rule, but here, of course, they suggest the wrong rule. The other option remains the access analyzer
upvoted 1 times

SkyZeroZx 1 year, 6 months ago
Selected Answer: B
keyword = AWS Identity and Access Management Access Analyzer
then B
upvoted 2 times

leehjworking 1 year, 7 months ago
Selected Answer: B
The code by God_is_love did not worked for me. I guess something has been changed.
The following code worked in my environment.
{
"source":["aws.access-analyzer"],
"detail-type":["Access Analyzer Finding"],
"detail":
{
"isPublic":[true]
}
}
upvoted 1 times

SkyZeroZx 1 year, 7 months ago
Selected Answer: B
Aws is letter B

Previous writing is a error
upvoted 1 times

SkyZeroZx 1 year, 7 months ago
Letter C
upvoted 1 times

SkyZeroZx 1 year, 7 months ago

Solution D will not meet the requirements because it will notify the data security team whenever an S3 bucket is not compliant with the cloudtrail-s3-dataevents-enabled rule, even if the bucket is not publicly exposed. The cloudtrail-s3-dataevents-enabled rule checks if at least one AWS CloudTrail trail is logging Amazon Simple Storage Service (Amazon S3) data events for all S3 buckets. If a bucket is not compliant with this rule, it does not mean that the bucket is publicly exposed. The bucket may simply not be logging S3 data events.
upvoted 2 times

SkyZeroZx 1 year, 7 months ago
Here are some reasons why an S3 bucket may not be logging S3 data events:

The bucket may not have a CloudTrail trail associated with it.
The CloudTrail trail for the bucket may not be enabled.
The CloudTrail trail for the bucket may not be configured to log S3 data events.
If the data security team is only interested in being notified when an S3 bucket becomes publicly exposed, then solution D is not the best solution. Solution B is a better solution because it will only notify the data security team when an S3 bucket becomes publicly exposed.
upvoted 1 times

y0eri 1 year, 7 months ago
Selected Answer: B
https://docs.aws.amazon.com/IAM/latest/UserGuide/access-analyzer-eventbridge.html
upvoted 1 times

mfsec 1 year, 9 months ago
Selected Answer: B
B eventbirdge and access analyser
upvoted 2 times

c73bf38 1 year, 10 months ago
Selected Answer: B
B is the correct solution because it uses AWS Identity and Access Management Access Analyzer to continuously monitor access control configurations and detect whether any S3 buckets have been configured to be publicly accessible. When a publicly accessible bucket is detected, an Amazon EventBridge rule is triggered, and the SNS topic is notified with the finding.
upvoted 7 times

masssa 1 year, 10 months ago
Selected Answer: B
Access Analyzer is to assess the access policy.
https://docs.aws.amazon.com/ja_jp/AmazonS3/latest/userguide/access-control-block-public-access.html
upvoted 2 times

[Removed] 1 year, 10 months ago
Selected Answer: B
https://aws.amazon.com/blogs/security/how-to-use-aws-iam-access-analyzer-api-to-automate-detection-of-public-access-to-aws-kms-keys/
upvoted 2 times

mdijoux25 1 year, 10 months ago
Selected Answer: B
https://docs.aws.amazon.com/AmazonS3/latest/userguide/access-analyzer.html
upvoted 2 times

spd 1 year, 10 months ago
Selected Answer: D
D by elimination rule
upvoted 2 times

Jay_2pt0_1 1 year, 7 months ago
I thought D, as well, but it seems everyone else things Access Analyzer.
upvoted 1 times

Question #174 *Topic 1*

A solutions architect needs to assess a newly acquired company's portfolio of applications and databases. The solutions architect must create a business case to migrate the portfolio to AWS. The newly acquired company runs applications in an on-premises data center. The data center is not well documented. The solutions architect cannot immediately determine how many applications and databases exist. Traffic for the applications is variable. Some applications are batch processes that run at the end of each month.

The solutions architect must gain a better understanding of the portfolio before a migration to AWS can begin.

Which solution will meet these requirements?

A. Use AWS Server Migration Service (AWS SMS) and AWS Database Migration Service (AWS DMS) to evaluate migration. Use AWS Service Catalog to understand application and database dependencies.

B. Use AWS Application Migration Service. Run agents on the on-premises infrastructure. Manage the agents by using AWS Migration Hub. Use AWS Storage Gateway to assess local storage needs and database dependencies.

C. Use Migration Evaluator to generate a list of servers. Build a report for a business case. Use AWS Migration Hub to view the portfolio. Use AWS Application Discovery Service to gain an understanding of application dependencies.

D. Use AWS Control Tower in the destination account to generate an application portfolio. Use AWS Server Migration Service (AWS SMS) to generate deeper reports and a business case. Use a landing zone for core accounts and resources.

Suggested Answer: *B*

Community vote distribution

C (97%)

spd Highly Voted 2 years, 4 months ago

Selected Answer: C

First need to evaluate

upvoted 16 times

c73bf38 Highly Voted 2 years, 4 months ago

Selected Answer: C

C. Use Migration Evaluator to generate a list of servers. Build a report for a business case. Use AWS Migration Hub to view the portfolio. Use AWS Application Discovery Service to gain an understanding of application dependencies.

upvoted 7 times

career360guru Most Recent 1 year, 7 months ago

Selected Answer: C

Option C

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: C

C for sure

upvoted 1 times

Roontha 2 years, 1 month ago

Answer : C

https://aws.amazon.com/migration-evaluator/

upvoted 2 times

F_Eldin 2 years, 1 month ago

Selected Answer: B

The emphasis is on applications. "Some applications are batch processes that run at the end of each month"

I do not understand why C is better than B

upvoted 1 times

helloworldabc 10 months, 1 week ago

just C

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: C

Use migration evaluator

upvoted 3 times

Question #175 *Topic 1*

A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.

Which solution will meet these requirements while providing the FASTEST storage performance?

A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster. Configure the ReplicaSet to mount the file system. Direct the application to store files in the file system. Configure AWS Backup to back up and retain copies of the data for 1 year.

B. Create an Amazon Elastic Block Store (Amazon EBS) volume. Enable the EBS Multi-Attach feature. Configure the ReplicaSet to mount the EBS volume. Direct the application to store files in the EBS volume. Configure AWS Backup to back up and retain copies of the data for 1 year.

C. Create an Amazon S3 bucket. Configure the ReplicaSet to mount the S3 bucket. Direct the application to store files in the S3 bucket. Configure S3 Versioning to retain copies of the data. Configure an S3 Lifecycle policy to delete objects after 1 year.

D. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locally. Use a third-party tool to back up the EKS cluster for 1 year.

Suggested Answer: *A*

Community vote distribution

A (100%)

c73bf38 Highly Voted 2 years, 4 months ago

Selected Answer: A

Explanation: Amazon EFS provides shared file storage that is highly available and durable. It is an ideal solution to share files between containers running on multiple instances in a cluster. Mounting an Amazon EFS file system on each subnet provides a shared file system for multiple instances running in different Availability Zones. Additionally, AWS Backup provides automated backup and recovery of Amazon EFS file systems.

upvoted 11 times

spd Highly Voted 2 years, 4 months ago

Selected Answer: A

EFS = Fastest storage performance compare to S3/EBS

upvoted 7 times

masssa 2 years, 4 months ago

I vote B.

I think EBS is faster than S3/EBS.

https://www.msp360.com/resources/blog/amazon-s3-vs-ebs-vs-efs/

upvoted 1 times

helloworldabc 10 months, 1 week ago

just A

upvoted 1 times

masssa 2 years, 4 months ago

typo.

EBS faster than S3/EFS.

upvoted 2 times

Musk 2 years, 4 months ago

I just read the question refers to multiple AZs, so B is not an option.

upvoted 9 times

AWSum1 8 months, 3 weeks ago

I missed this too good spot

upvoted 1 times

career360guru Most Recent 1 year, 7 months ago

Selected Answer: A

Option A

upvoted 1 times

joleneinthebackyard 1 year, 8 months ago

Selected Answer: A

A: sounds valid

B: EBS multi attach can only do same AZ -> out

C: S3 is for durability, not for perfomance

D: can drop when seeing third party tool.

upvoted 5 times

NikkyDicky 1 year, 12 months ago

Selected Answer: A

A - EFS for multi-AZ

upvoted 2 times

dkx 1 year, 12 months ago

A. Yes, because Amazon EFS offers you the choice of creating file systems using Standard or One Zone storage classes. Standard storage classes store data with and across multiple AZs.
https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/run-stateful-workloads-with-persistent-data-storage-by-using-amazon-efs-on-amazon-eks-with-aws-fargate.html

B. No, because Amazon EBS Multi-Attach enabled volumes can be attached to up to 16 Linux instances built on the Nitro System that are in the same Availability Zone. We need to solve for "nodes in multiple Availability Zones"
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html

C. No, because if you're looking to run file-based applications that need to collaborate or coordinate on shared data across instances or users, AWS recommends fully managed file services, such as Amazon FSx or Amazon Elastic File System (EFS).

D. No, because the company needs to back up the files, not backup the EKS Cluster.

upvoted 4 times

mfsec 2 years, 3 months ago

Selected Answer: A

A for sure

upvoted 2 times

ramyaram 2 years, 3 months ago

Selected Answer: A

Keyword here is multiple small files and shared between multiple clusters

upvoted 3 times

God_Is_Love 2 years, 3 months ago

Selected Answer: A

In the past, EBS can be attached only to one ec2 instance but not anymore but there are limitations like - it works only on io1/io2 instance types and many others as described here. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html
EFS has shareable storage

In terms of performance, Amazon EFS is optimized for workloads that require high levels of aggregate throughput and IOPS, whereas EBS is optimized for low-latency, random access I/O operations. Amazon EFS is designed to scale throughput and capacity automatically as your storage needs grow, while EBS volumes can be resized on demand.

upvoted 3 times

Zek 2 years, 3 months ago

I support A since their is a multi-AZ requirement.

https://repost.aws/questions/QUK2RANw1QTKCwpDUwCCI72A/efs-vs-ebs-mult-attach

EFS is also designed for high availability and high durability. To achieve these levels of availability and durability, EFS automatically replicates data within and across 3 Availability Zones, with no single points of failure. EBS multi-attach volumes can be used for clients within a single Availability Zone.

upvoted 1 times

Sarutobi 2 years, 3 months ago

Selected Answer: A

When you have an EKS cluster and use the EBS that is local to the node, only Pods running on that node have access to the storage. If the node starts on any other Pod, it will potentially break. There are ways to fix this, but they are beyond this question. I believe we need shared fast storage here, so it should be S3 vs EFS the decision.

upvoted 3 times

Musk 2 years, 4 months ago

I've been reding here and there, and B does not seem that feasible, although if supported it would be faster than A.

upvoted 2 times

Question #176 *Topic 1*

A company runs a customer service center that accepts calls and automatically sends all customers a managed, interactive, two-way experience survey by text message. The applications that support the customer service center run on machines that the company hosts in an on-premises data center. The hardware that the company uses is old, and the company is experiencing downtime with the system. The company wants to migrate the system to AWS to improve reliability.

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

A. Use Amazon Connect to replace the old call center hardware. Use Amazon Pinpoint to send text message surveys to customers.

B. Use Amazon Connect to replace the old call center hardware. Use Amazon Simple Notification Service (Amazon SNS) to send text message surveys to customers.

C. Migrate the call center software to Amazon EC2 instances that are in an Auto Scaling group. Use the EC2 instances to send text message surveys to customers.

D. Use Amazon Pinpoint to replace the old call center hardware and to send text message surveys to customers.

Suggested Answer: *A*

Community vote distribution

A (100%)

God_Is_Love Highly Voted 1 year, 9 months ago

Selected Answer: A

Amazon Connect is a cloud-based contact center service that allows you to set up a virtual call center for your business. It provides an easy-to-use interface for managing customer interactions through voice and chat. Amazon Connect integrates with other AWS services, such as Amazon S3 and Amazon Kinesis, to help you collect, store, and analyze customer data for insights into customer behavior and trends.

On the other hand, Amazon Pinpoint is a marketing automation and analytics service that allows you to engage with your customers across different channels, such as email, SMS, push notifications, and voice. It helps you create personalized campaigns based on user behavior and enables you to track user engagement and retention.

While both services allow you to communicate with your customers, they serve different purposes. Amazon Connect is focused on customer support and service, while Amazon Pinpoint is focused on marketing and engagement.

upvoted 13 times

pichunya Most Recent 1 month ago

Selected Answer:

amazon pinpoint EoS 2025/5/20

upvoted 1 times

alexsanteeno 1 year ago

"LEAST OPERATIONAL OVERHEAD" - is key word in a question. Its not so easy to migrate any on-premise infra to any AWS. Looking at the answers here I see no one eve done that before and just answering as from AWS docs.
The easiest way to migrate any on-premise infra - ec2

upvoted 2 times

career360guru 1 year, 1 month ago

Selected Answer: A

Option A

upvoted 1 times

rrrrrrrrr1 1 year, 5 months ago

Why not b though? SNS is easy as heck to use.

upvoted 1 times

VerRi 10 months, 1 week ago

"managed, interactive, two-way experience" means a personalised and customised message, so it should be Pinpoint here.

upvoted 5 times

rrrrrrrrr1 1 year, 5 months ago

nvm text message surveys are probably a pinpoint thing. I was thinking like a link to a survey.

upvoted 3 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

A - basic AWS connect use case

upvoted 1 times

Maria2023 1 year, 6 months ago

Selected Answer: A

Amazon connect + Pinpoint are the best choice here

upvoted 1 times

Roontha 1 year, 7 months ago

Answer: A

upvoted 1 times

mfsec 1 year, 9 months ago

Selected Answer: A

Use Amazon Connect to replace the old call center hardware. Use Amazon Pinpoint to send text message surveys to customers.

upvoted 1 times

c73bf38 1 year, 10 months ago

Selected Answer: A

The solution that will meet the company's requirements with the LEAST ongoing operational overhead and send two-way experience survey is to use Amazon Connect to replace the old call center hardware and use Amazon Pinpoint to send text message surveys to customers. Amazon Connect is a fully managed, cloud-based contact center service that is easy to set up and configure, while Amazon Pinpoint can be used to send text message surveys and gather responses. By using these services, the company can offload the operational overhead of running and maintaining the call center hardware and survey system to AWS.

upvoted 4 times

spd 1 year, 10 months ago

Selected Answer: A

https://docs.aws.amazon.com/pinpoint/latest/userguide/channels-sms-two-way.html

upvoted 2 times

Question #177 *Topic 1*

A company is building a call center by using Amazon Connect. The company's operations team is defining a disaster recovery (DR) strategy across AWS Regions. The contact center has dozens of contact flows, hundreds of users, and dozens of claimed phone numbers.

Which solution will provide DR with the LOWEST RTO?

A. Create an AWS Lambda function to check the availability of the Amazon Connect instance and to send a notification to the operations team in case of unavailability. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes. After notification, instruct the operations team to use the AWS Management Console to provision a new Amazon Connect instance in a second Region. Deploy the contact flows, users, and claimed phone numbers by using an AWS CloudFormation template.

B. Provision a new Amazon Connect instance with all existing users in a second Region. Create an AWS Lambda function to check the availability of the Amazon Connect instance. Create an Amazon EventBridge rule to invoke the Lambda function every 5 minutes. In the event of an issue, configure the Lambda function to deploy an AWS CloudFormation template that provisions contact flows and claimed numbers in the second Region.

C. Provision a new Amazon Connect instance with all existing contact flows and claimed phone numbers in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance. Create an Amazon CloudWatch alarm for failed health checks. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions all users. Configure the alarm to invoke the Lambda function.

D. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region. Create an Amazon Route 53 health check for the URL of the Amazon Connect instance. Create an Amazon CloudWatch alarm for failed health checks. Create an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. Configure the alarm to invoke the Lambda function.

Suggested Answer: *D*

Community vote distribution

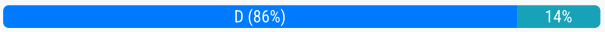

nyxs_19 Highly Voted 2 years, 4 months ago

Selected Answer: D

The solution that will provide DR with the LOWEST RTO (Recovery Time Objective) is option D.

Option D provisions a new Amazon Connect instance with all existing users and contact flows in a second Region. It also sets up an Amazon Route 53 health check for the URL of the Amazon Connect instance, an Amazon CloudWatch alarm for failed health checks, and an AWS Lambda function to deploy an AWS CloudFormation template that provisions claimed phone numbers. This option allows for the fastest recovery time because all the necessary components are already provisioned and ready to go in the second Region. In the event of a disaster, the failed health check will trigger the AWS Lambda function to deploy the CloudFormation template to provision the claimed phone numbers, which is the only missing component.

upvoted 10 times

spd Highly Voted 2 years, 4 months ago

Selected Answer: D

D looks most appropriate

upvoted 9 times

29fb203 Most Recent 3 months, 1 week ago

Selected Answer: C

C accounts for the phone numbers and all other resources. D doesn't

upvoted 1 times

Sin_Dan 8 months, 1 week ago

Selected Answer: C

Setting up phone numbers is more complex and time consuming, than setting up users. Option D waits until the disaster happens to provision the phone numbers. Option C is right, because it is quicker as compared to option D. Also, it makes sure the users are not duplicated upfront.

upvoted 1 times

cashyc 8 months, 2 weeks ago

Selected Answer: C

by pre-provisioning a new Amazon Connect instance in a second AWS Region with the necessary contact flows and phone numbers already in place. The remaining task at the time of disaster recovery is to deploy the users, which can be done using an AWS Lambda function triggered by a CloudWatch alarm when the primary instance becomes unavailable, as determined by a Route 53 health check.

upvoted 1 times

marszalekm 1 year, 5 months ago

Amazon Connect is not on the list of services required for this exam. At least as of 08.01.24 https://d1.awsstatic.com/training-and-certification/docs-sa-pro/AWS-Certified-Solutions-Architect-Professional_Exam-Guide.pdf

upvoted 6 times

career360guru 1 year, 7 months ago

Selected Answer: D

Option D

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: D

Amazon Connect gives you a URL, for which you can add a record in route 53 and hence have a health check.

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: D

D seems to fit all requirements, however C & D seem to be very similar. Only difference is whether to upload users or phone numbers through Cloud Formation. It seems users, routing profiles, queues, and flows get created with ReplicateInstance API https://docs.aws.amazon.com/connect/latest/adminguide/create-replica-connect-instance.html

upvoted 3 times

MRL110 1 year, 11 months ago

Selected Answer: B

Apparently Route 53 can't manage Amazon Connect DNS names or health checks. https://docs.aws.amazon.com/connect/latest/adminguide/update-your-connect-domain.html#new-domain-custom

upvoted 1 times

NikkyDicky 1 year, 12 months ago

Selected Answer: D

D i guess

upvoted 1 times

Maria2023 2 years ago

Selected Answer: B

I vote for B since I was not able to find a way to make Route53 serve the Amazon connect URL and therefore it cannot perform healthcheck. If someone has more information on this - please share

upvoted 1 times

SkyZeroZx 2 years, 1 month ago

why not letter C

"CloudFormation template that provisions all users" insted of "CloudFormation template that provisions claimed phone numbers" of letter D

upvoted 3 times

dev112233xx 2 years, 2 months ago

Selected Answer: B

I'm voting B because i don't think it's possible to use Amazon Route 53 health check to verify the availability of Amazon Connect

upvoted 1 times

Eshu2009 2 years, 3 months ago

why not C?

upvoted 1 times

ninomfr64 1 year, 5 months ago

I think, but I was not able to very it, that if your instance is active and you have phone numbers configured it is receiving actual phone traffic that is a and Active/Active scenario, however you do not have users (aka Agents) configured to handle calls. This is just me guessing

upvoted 3 times

shmoeee 4 months, 2 weeks ago

Same thinking i had

upvoted 1 times

mfsec 2 years, 3 months ago

Selected Answer: D

D. Provision a new Amazon Connect instance with all existing users and contact flows in a second Region.

upvoted 3 times

c73bf38 2 years, 4 months ago

Selected Answer: D

D is the better solution.

upvoted 3 times

Question #178 *Topic 1*

A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs ETL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.

The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data. The customers also need access to the most recent data when the company publishes the data.

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

A. Use AWS Data Exchange for APIs to share data with customers. Configure subscription verification. In the AWS account of the company that produces the data, create an Amazon API Gateway Data API service integration with Amazon Redshift. Require the data customers to subscribe to the data product.

B. In the AWS account of the company that produces the data, create an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster. Configure subscription verification. Require the data customers to subscribe to the data product.

C. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically. Use AWS Data Exchange for S3 to share data with customers. Configure subscription verification. Require the data customers to subscribe to the data product.

D. Publish the Amazon Redshift data to an Open Data on AWS Data Exchange. Require the customers to subscribe to the data product in AWS Data Exchange. In the AWS account of the company that produces the data, attach IAM resource-based policies to the Amazon Redshift tables to allow access only to verified AWS accounts.

Suggested Answer: *B*

Community vote distribution

youngmanaws Highly Voted 1 year, 2 months ago

Selected Answer: B

The company wants to confirm the identities of the customers before the company shares data. The customers also need access to the most recent data when the company publishes the data. With B, customer can get data from Redshift directly with no time lag and additional operations.

upvoted 11 times

renegadedme Highly Voted 1 year, 2 months ago

Selected Answer: B

I think it's B.

According to https://aws.amazon.com/data-exchange/why-aws-data-exchange/redshift-data-tables/

Customers can find and subscribe to third-party data in AWS Data Exchange and directly query the data in minutes in Amazon Redshift without extracting, transforming, or loading it.

In B, customers can query Redshift directly. No need to use S3 periodically. Minimizes operational overhead.

upvoted 9 times

NikkyDicky Most Recent 11 months, 4 weeks ago

Selected Answer: B

it's a B

upvoted 1 times

SmileyCloud 12 months ago

Selected Answer: B

Keyword is datashare

https://docs.aws.amazon.com/redshift/latest/dg/adx-getting-started.html

upvoted 5 times

easytoo 1 year ago

b-b-b-b-bb-b-b-b-b-b

LEAST operational overhead...

Option (A) uses AWS Data Exchange for APIs, which requires you to create an Amazon API Gateway Data API service integration with Amazon Redshift. This is a more complex solution than using a datashare.

Option (C) uses AWS Data Exchange for S3, which requires you to download the data from Amazon Redshift to Amazon S3 periodically. This is also a more complex solution than using a datashare.

Option (D) publishes the data to an Open Data on AWS Data Exchange, which does not allow you to configure subscription verification. This means that anyone can access the data, which is not ideal for a company that wants to protect its proprietary algorithms.

upvoted 3 times

TECHNOWARRIOR 1 year ago

AWS Data Exchange for APIs enables customers to discover and utilize third-party APIs in the cloud, with authentication using AWS IAM credentials and SDKs. It simplifies access permissions and governance. Users can access data APIs from numerous providers. On the other hand, AWS Data Exchange Datashare focuses on licensing access to Amazon Redshift data. It utilizes AWS-native authentication and automatically adds customers as data consumers. With read-only access, customers can retrieve objects from datashares. While both services integrate with AWS, Data Exchange for APIs is geared towards API usage, while Data Exchange Datashare is centered around licensing access to Amazon Redshift data.

upvoted 5 times

Roontha 1 year, 1 month ago

Answer : B

https://www.youtube.com/watch?v=BeIoTSql4IM

(AWS Data Exchange for Amazon Redshift demo | Amazon Web Services)

upvoted 3 times

Sarutobi 1 year, 1 month ago

Selected Answer: B

B is the closest one but is not correct either.

https://docs.amazonaws.cn/en_us/redshift/latest/dg/adx-getting-started-producer.html, like every thing else in AWS you need policy to grant access and that is missing in B.

upvoted 2 times

nqg54118 1 year, 2 months ago

Selected Answer: C

□□□□□□□□□□□□□□□□□□□□S3

upvoted 1 times

easytoo 1 year ago

yup! was about to say the same.

upvoted 6 times

yorkicurke 7 months, 3 weeks ago

hahahaaha

upvoted 1 times

OCHT 1 year, 2 months ago

Selected Answer: C

The correct answer is C. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically. Use AWS Data Exchange for S3 to share data with customers. Configure subscription verification. Require the data customers to subscribe to the data product.

Exporting the data to an Amazon S3 bucket periodically ensures that customers have access to the most recent data when the company publishes it.
AWS Data Exchange for S3 allows you to share data with customers easily and manage their subscriptions.
Subscription verification helps confirm the identity of customers before sharing data with them.
This solution minimizes operational overhead as it leverages AWS Data Exchange and Amazon S3, which are managed services.
The unique keywords combination in this option that makes it easier to remember is Amazon S3, AWS Data Exchange, and subscription verification.

upvoted 2 times

Yowie351 1 year, 2 months ago

Selected Answer: B

Answer is B. https://aws.amazon.com/data-exchange/?adx-cards2.sort-by=item.additionalFields.eventDate&adx-cards2.sort-order=desc

upvoted 2 times

Question #179 *Topic 1*

A solutions architect is designing a solution to process events. The solution must have the ability to scale in and out based on the number of events that the solution receives. If a processing error occurs, the event must move into a separate queue for review.

Which solution will meet these requirements?

A. Send event details to an Amazon Simple Notification Service (Amazon SNS) topic. Configure an AWS Lambda function as a subscriber to the SNS topic to process the events. Add an on-failure destination to the function. Set an Amazon Simple Queue Service (Amazon SQS) queue as the target.

B. Publish events to an Amazon Simple Queue Service (Amazon SQS) queue. Create an Amazon EC2 Auto Scaling group. Configure the Auto Scaling group to scale in and out based on the ApproximateAgeOfOldestMessage metric of the queue. Configure the application to write failed messages to a dead-letter queue.

C. Write events to an Amazon DynamoDB table. Configure a DynamoDB stream for the table. Configure the stream to invoke an AWS Lambda function. Configure the Lambda function to process the events.

D. Publish events to an Amazon EventBndge event bus. Create and run an application on an Amazon EC2 instance with an Auto Scaling group that is behind an Application Load Balancer (ALB). Set the ALB as the event bus target. Configure the event bus to retry events. Write messages to a dead-letter queue if the application cannot process the messages.

Suggested Answer: *B*

Community vote distribution

A (62%) B (38%)

Sarutobi Highly Voted 2 years, 2 months ago

Selected Answer: B

I would go with B just because of the wording. I believe A should work just fine, but the question asks for "scale in and out based on the number of events." In my opinion, that is what SNS->Lambda->SQS(DLQ) would do, too; I think the SNS->Lambda scale in/out behavior is more implicit. So I will go with B here because it is more explicit.

upvoted 32 times

SuperDuperPooperScooper Highly Voted 1 year, 7 months ago

Selected Answer: A

Configuring scaling based on the age of the oldest message is nowhere near as good as scaling based on size of the Queue for this use case.

age of the oldest message will grow linearly based on time. If there is a dramatic spike in the Queue size due to increased traffic, like 100X increase in size. Then the queue will have grown a lot but the oldest message will only increase in age linearly, so the scaling will not be able to realize how much the workload has increased.

upvoted 11 times

helloworldabc 10 months, 1 week ago

just B

upvoted 1 times

mns0173 10 months, 4 weeks ago

there will just be a lag in scaling, but eventually this metric will scale as needed

upvoted 1 times

sonyaws 1 year, 7 months ago

makes sense

upvoted 1 times

jainparag1 1 year, 7 months ago

very good explanation. Moreover, go serverless as much as possible. EC2 vs Lambda - Lamda is always preferred.

upvoted 1 times

Kaps443 Most Recent 2 weeks, 5 days ago

Selected Answer: B

This one is easy haven't people heard about SQS dead-letter queues (DLQs).

upvoted 1 times

eesa 1 month, 2 weeks ago

Selected Answer: B

B is correct:

Scalability: The EC2 Auto Scaling group can automatically scale in and out based on the SQS metric (ApproximateAgeOfOldestMessage), which reflects how long messages have been waiting to be processed.

Error handling: SQS supports dead-letter queues (DLQs) to isolate and handle failed messages for later analysis or reprocessing.

Decoupled architecture: SQS enables a loosely coupled and fault-tolerant system design.

upvoted 1 times

f3f4935 1 month, 4 weeks ago

Selected Answer: B

also think B will be good there

upvoted 1 times

CAIYasia 2 months ago

Selected Answer: B

I would go B for the DLQ

upvoted 1 times

itsjunukim 4 months ago

Selected Answer: B

By utilizing the ApproximateAgeOfOldestMessage metric, you can scale out and scale in based on the workload, ensuring that your application can handle increases in traffic.

upvoted 1 times

820b83f 4 months, 2 weeks ago

Selected Answer: B

100 % its B, My reasons are:

1. SNS is for pub/sub, not event processing. SNS sends events to multiple subscribers but does not provide queue-based scaling.

2. Lambda also has concurrency limits, which might cause failures at high event rates.

upvoted 1 times

kylix75 5 months, 1 week ago

Selected Answer: B

The correct answer is B.

Reasons:
1. SQS + Auto Scaling provides event-based scalability
2. ApproximateAgeOfOldestMessage metric enables workload-based scaling
3. SQS native dead-letter queue handles error messages
4. Most resilient and cost-effective solution for event processing at scale

Issues with other options:
A: SNS doesn't store messages for reprocessing
C: DynamoDB Streams has scalability and retention limitations
D: ALB + EC2 is more complex and expensive than serverless processing

upvoted 1 times

ahhatem 6 months, 2 weeks ago

Selected Answer: B

While A would probably work fine most of the time, B is more resilient. Once a message is in the Q, it will either be marked as complete or go to DLQ. In A, in edge cases like lambda exceeding concurrency limit, the message would be throttled after the SNS returns success to the sender.... Without SNS DLQ, the message would be lost.

upvoted 3 times

FZA24 8 months ago

Selected Answer: B

scale in scale out => ALB

a separate queue => DLQ

upvoted 3 times

Woody1848 8 months, 1 week ago

Selected Answer: A

• By sending event details to an Amazon SNS topic and configuring an AWS Lambda function as a subscriber, the solution automatically scales with the number of incoming events.

• Lambda functions scale in and out based on the event load without manual intervention.

• Adding an on-failure destination to the Lambda function that targets an Amazon SQS queue ensures that any processing errors move the event into a separate queue for review.

• This setup meets both the scalability and error-handling requirements efficiently.

upvoted 2 times

Sin_Dan 8 months, 1 week ago

Selected Answer: B

Option A uses Lambda to process the solution. However, we don't know if the processing finishes within 15 mins or not. Also, SNS isn't as well-suited for handling large event queues as SQS, and scaling based on message queue metrics is not supported in this configuration.

So, the correct option is definitely B.

upvoted 2 times

Daniel76 8 months, 3 weeks ago

Selected Answer: B

Only B and D mention about reviewing error in a separate queue by dead letter Q, with D never use SQS where this is supported.

upvoted 2 times

Syre 10 months, 2 weeks ago

Selected Answer: B

People who are choosing A. Have you done associate level certs as this does not make any sense. You shouldnt be attempting this exam if that's how lost you are.

upvoted 2 times

amsf96 7 months, 1 week ago

chill bro

upvoted 5 times

ChungFTF 10 months, 3 weeks ago

Selected Answer: B

The Auto Scaling group of EC2 instances can automatically adjust the number of instances based on the ApproximateAgeOfOldestMessage metric. This ensures that the solution scales dynamically with the volume of events, maintaining efficient processing.

https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-using-sqs-queue.html

upvoted 2 times

Linuslin 10 months, 2 weeks ago

Dude, check you link again, which says "ApproximateNumberOfMessages" not " ApproximateAgeOfOldestMessage", so answer will be option A.

upvoted 3 times

tsangckl 11 months, 1 week ago

Selected Answer: B

B for sure

upvoted 2 times

Question #180 *Topic 1*

A company runs a processing engine in the AWS Cloud. The engine processes environmental data from logistics centers to calculate a sustainability index. The company has millions of devices in logistics centers that are spread across Europe. The devices send information to the processing engine through a RESTful API.

The API experiences unpredictable bursts of traffic. The company must implement a solution to process all data that the devices send to the processing engine. Data loss is unacceptable.

Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) for the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create a listener and a target group for the ALB Add the SQS queue as the target. Use a container that runs in Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type to process messages in the queue.

B. Create an Amazon API Gateway HTTP API that implements the RESTful API. Create an Amazon Simple Queue Service (Amazon SQS) queue. Create an API Gateway service integration with the SQS queue. Create an AWS Lambda function to process messages in the SQS queue.

C. Create an Amazon API Gateway REST API that implements the RESTful API. Create a fleet of Amazon EC2 instances in an Auto Scaling group. Create an API Gateway Auto Scaling group proxy integration. Use the EC2 instances to process incoming data.

D. Create an Amazon CloudFront distribution for the RESTful API. Create a data stream in Amazon Kinesis Data Streams. Set the data stream as the origin for the distribution. Create an AWS Lambda function to consume and process data in the data stream.

Suggested Answer: *B*

Community vote distribution

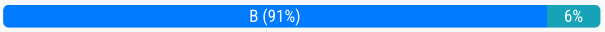

momo3321 Highly Voted 2 years, 1 month ago

Selected Answer: B

Option A is incorrect because Application Load Balancer (ALB) can't directly target an Amazon SQS queue.

Option C is incorrect because while Amazon API Gateway and EC2 Auto Scaling can handle high loads, they don't provide a built-in mechanism to ensure that all messages are processed without loss.

Option D is incorrect because Amazon CloudFront is a content delivery network (CDN), and it is not typically used to handle incoming API requests. It is primarily used to cache and deliver content to users.

upvoted 20 times

bjexamprep Highly Voted 1 year, 6 months ago

Selected Answer: B

In real life, I wouldn't trust SQS to handle such large amount of data.

upvoted 6 times

altonh Most Recent 4 months ago

Selected Answer: D

A - SQS as an ALB target is wrong

B - Cannot integrate an AWS service using an HTTP API Gateway

C - Data cannot be passed to EC2 because the integrated AWS service is an ASG.

upvoted 1 times

vip2 11 months, 3 weeks ago

Selected Answer: B

Restful API is not REST API, so HTTP API-GW + SQS

upvoted 2 times

jAtlas7 7 months ago

REST APIs and HTTP APIs are both RESTful API products. Ref: https://docs.aws.amazon.com/apigateway/latest/developerguide/http-api-vs-rest.html

upvoted 1 times

nzin4x 1 year, 4 months ago

but normally API gateway can not handle high burst request. it will make 429 too many requests error.

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: B

yes you can integrate API Gateway HTTP Api with SQS

upvoted 2 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: B

KDS need to implement Sharding for unpredictable bursts

upvoted 1 times

rxhan 1 year, 11 months ago

Similar to #179

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

B is right

upvoted 1 times

Roontha 2 years, 1 month ago

Answer : B

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: B

I agree with B

https://aws.amazon.com/blogs/architecture/things-to-consider-when-you-build-rest-apis-with-amazon-api-gateway/

This pattern can decouple the data ingestion from the data processing.

"you should look for opportunities to design an asynchronous, loosely coupled architecture. A decoupled architecture separates the data ingestion from the data processing and allows you to scale each system separately"

upvoted 2 times

AMEJack 2 years, 1 month ago

Selected Answer: B

Kinesis DataStreams can't be the origin for the CloudFront

upvoted 2 times

mrfretz 2 years, 2 months ago

Selected Answer: D

Kinesis retention

upvoted 1 times

mrfretz 2 years, 2 months ago

Selected Answer: B

Kinesis retention

upvoted 1 times

mrfretz 2 years, 2 months ago

Answer D, sorry typo

upvoted 1 times

Sarutobi 2 years, 2 months ago

B is the best option.

upvoted 1 times

Question #181 *Topic 1*

A company is designing its network configuration in the AWS Cloud. The company uses AWS Organizations to manage a multi-account setup. The company has three OUs. Each OU contains more than 100 AWS accounts. Each account has a single VPC, and all the VPCs in each OU are in the same AWS Region.

The CIDR ranges for all the AWS accounts do not overlap. The company needs to implement a solution in which VPCs in the same OU can communicate with each other but cannot communicate with VPCs in other OUs.

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

A. Create an AWS CloudFormation stack set that establishes VPC peering between accounts in each OU. Provision the stack set in each OU.

B. In each OU, create a dedicated networking account that has a single VPC. Share this VPC with all the other accounts in the OU by using AWS Resource Access Manager (AWS RAM). Create a VPC peering connection between the networking account and each account in the OU.

C. Provision a transit gateway in an account in each OU. Share the transit gateway across the organization by using AWS Resource Access Manager (AWS RAM). Create transit gateway VPC attachments for each VPC.

D. In each OU, create a dedicated networking account that has a single VPC. Establish a VPN connection between the networking account and the other accounts in the OU. Use third-party routing software to route transitive traffic between the VPCs.

Suggested Answer: *D*

Community vote distribution

SK_Tyagi Highly Voted 1 year, 10 months ago

Selected Answer: C

Fits the use case

https://aws.amazon.com/transit-gateway/

upvoted 13 times

SK_Tyagi 1 year, 10 months ago

https://docs.aws.amazon.com/vpc/latest/tgw/transit-gateway-isolated.html

upvoted 2 times

ninomfr64 Highly Voted 1 year, 5 months ago

Option C is very poorly worded: "Provision a transit gateway in an account in each OU" to me this results in having 3 Transit Gateways, but then it go ahead just referring to a single Transit Gateway "Share the transit gateway across the organization ..."

upvoted 7 times

bhanus Most Recent 6 months ago

Selected Answer: C

TGW would be used to create hub and spoke. VPCs are in same region so tgw can be shared via RAM.

Answer is C

upvoted 1 times

43c89f4 1 year, 2 months ago

Typical transit gateway use case

upvoted 1 times

bjexamprep 1 year, 3 months ago

Selected Answer: C

The question is asking “a solution in which VPCs in the same OU can communicate with each other but cannot communicate with VPCs in other OUs”.

A: it works. But it may create 2500+ VPC peering in each OU

B: It works. But it may create 2500+ VPC peering in each OU

C: This is wrong, cause it is sharing the transit gateway to all the account in the organization instead of sharing to all the account in that OU.

D: That means 2500+ VPN connections in each OU and cost a lot of internet bandwidth.

I guess the C was worded with mistake. It should be sharing the transit gateway to the accounts in each OU and create VPC attachment for each VPC in that OU.

upvoted 5 times

Sin_Dan 8 months, 1 week ago

I don't understand why there are so many poorly written questions and options in the AWS exams. I am wondering if we are writing an exam for English or AWS. Many questions are just elongated for adding complexity. Not a right way to assess technical skills of a person based on their English skills.

upvoted 2 times

VerRi 1 year, 4 months ago

Selected Answer: A

The requirement said, "VPCs in the same OU can communicate with each other but cannot communicate with VPCs in other OUs". There is no reason to share the TGW across the organisation with RAM because it will enable cross OUs communication.

upvoted 1 times

itsjunukim 4 months ago

VPCs within the same OU can communicate with each other. Each OU has 100 accounts, and having all 100 accounts perform VPC peering would be inefficient.

upvoted 1 times

learnwithaniket 1 year, 6 months ago

Selected Answer: A

"Least operational overhead"

A is correct.

C creating Transit Gateway in each account.. and there are more than 100 accounts in each OU. Which is time consuming and requires lot of efforts.

upvoted 2 times

chicagobeef 1 year, 5 months ago

"A" would mean having 1:1 peering attachments with EACH ACCOUNT which is too much operational overhead. A transit gateway is more viable so it's "C".

upvoted 4 times

jainparag1 1 year, 7 months ago

Selected Answer: A

typical use case of intra region peering with transit gateway.

upvoted 1 times

jainparag1 1 year, 7 months ago

oops right answer is 'C'.

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C

upvoted 3 times

rlf 1 year, 8 months ago

C.

Transit gateway and RAM is a regional service.

AWS RAM is a Regional service, and a resource share is Regional. Therefore, a resource share can contain resources from the same AWS Region as the resource share, and any supported global resources.

https://docs.aws.amazon.com/ram/latest/userguide/working-with-regional-vs-global.html

https://docs.aws.amazon.com/ram/latest/userguide/getting-started-sharing.html#getting-started-sharing-orgs

upvoted 6 times

LuongTo 7 months ago

the best explanation, share across but the same Region -> same OU

upvoted 1 times

MRL110 1 year, 11 months ago

Selected Answer: A

A for two reasons:

1. Sharing the TGW with the entire organization (C) will make every VPC in every account propagate its subnet in the default TGW route table which will enable organization-wide communication which is categorically prohibited by the question.

2. The question only says more than 100 accounts and 1 VPC per account. It does not mention anything about 125+ VPCs. Plus the peerings are being created by stack sets so there's automation involved. So I believe A is the only solution here.

upvoted 1 times

MRL110 1 year, 11 months ago

Disabling default route table association/propagation could be a solution for TGW, but creating 100s of VPC attachments manually is too much operational overhead.

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: C

I thik C

upvoted 3 times

dkx 1 year, 12 months ago

C. Yes, because, Transit Gateway is a managed service from AWS that acts as a hub interconnecting VPCs and VPN connections within a single region. It allows you to build more complex networks without the need for VPC peering.
Similar to: https://aws.amazon.com/blogs/networking-and-content-delivery/automating-aws-transit-gateway-attachments-to-a-transit-gateway-in-a-central-account/

A,B. No, because a VPC peering connection has a limit of 125 Active VPC peering connections per VPC. In this case, each OU contains MORE THAN 100 AWS accounts -- this could mean 101 accounts or 10001 accounts.

D. No, because this is not the answer choice with the LEAST operational overhead. Third-party routing software is not required to route transitive traffic between the VPCs.

upvoted 5 times

xflare 1 year, 10 months ago

I believe in this context the organization is the OU, not the entire company. The company is referred to as "the company".
Therefore it's C.

upvoted 1 times

pupsik 2 years ago

Selected Answer: C

A separate transit GW for each OU.

upvoted 2 times

Maria2023 2 years ago

Selected Answer: C

The answer should be C. Since VPC peering is not transitive then for 100+ accounts in OU then we'll breach the limit of 125. As for VPN - I wouldn't use VPN to connect AWS resources - I don't know even if that's possible

upvoted 2 times

Jackhemo 2 years ago

Olabiba.ai says C.

upvoted 2 times

Ashas 2 years ago

I have an exam on 27th june, what question set should I prepare? I have only done from Question#1 to Question#181 yet. Please help

upvoted 2 times

Roontha 2 years, 1 month ago

Answer : C

Reference : https://catalog.workshops.aws/networking/en-US/intermediate/6-vpc-peering/10-vpc-peering-overview

upvoted 1 times

Question #182 *Topic 1*

A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large important documents within the application with the following requirements:

1. The data must be highly durable and available
2. The data must always be encrypted at rest and in transit
3. The encryption key must be managed by the company and rotated periodically

Which of the following solutions should the solutions architect recommend?

A. Deploy the storage gateway to AWS in file gateway mode. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.

B. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.

C. Use Amazon DynamoDB with SSL to connect to DynamoDB. Use an AWS KMS key to encrypt DynamoDB objects at rest.

D. Deploy instances with Amazon EBS volumes attached to store this data. Use EBS volume encryption using an AWS KMS key to encrypt the data.

Suggested Answer: *B*

Community vote distribution

B (100%)

SkyZeroZx Highly Voted 2 years ago

if you have come far it means that you are persistent, good luck in your exam

upvoted 33 times

kgpoj 10 months, 3 weeks ago

Man, what can I say

upvoted 2 times

easytoo 2 years ago

My man. Respect, we are all cloud brothers here.

upvoted 10 times

joleneinthebackyard 1 year, 8 months ago

I went backward, does it count?

upvoted 8 times

gutomarson Most Recent 12 months ago

Answer is B

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: B

Easy breezy

upvoted 2 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

its a b

upvoted 1 times

Maria2023 2 years ago

Selected Answer: B

At least an easy one - the provided configuration for S3 in B satisfies the requirements for encryption, durability and availability

upvoted 3 times

Alabi 2 years ago

Selected Answer: B

B for sure

upvoted 1 times

erhard 2 years ago

Not C because _large_ documents and
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ServiceQuotas.html#limits-items

upvoted 2 times

Alabi 2 years ago

Selected Answer: B

Definitely B

upvoted 2 times

kfrum4 2 years ago

Selected Answer: B

Answer: B

upvoted 1 times

AMEJack 2 years, 1 month ago

Selected Answer: B

Answer is B

upvoted 2 times

Roontha 2 years, 1 month ago

Answer : B

upvoted 2 times

Question #183 *Topic 1*

A company's public API runs as tasks on Amazon Elastic Container Service (Amazon ECS). The tasks run on AWS Fargate behind an Application Load Balancer (ALB) and are configured with Service Auto Scaling for the tasks based on CPU utilization. This service has been running well for several months.

Recently, API performance slowed down and made the application unusable. The company discovered that a significant number of SQL injection attacks had occurred against the API and that the API service had scaled to its maximum amount.

A solutions architect needs to implement a solution that prevents SQL injection attacks from reaching the ECS API service. The solution must allow legitimate traffic through and must maximize operational efficiency.

Which solution meets these requirements?

A. Create a new AWS WAF web ACL to monitor the HTTP requests and HTTPS requests that are forwarded to the ALB in front of the ECS tasks.

B. Create a new AWS WAF Bot Control implementation. Add a rule in the AWS WAF Bot Control managed rule group to monitor traffic and allow only legitimate traffic to the ALB in front of the ECS tasks.

C. Create a new AWS WAF web ACL. Add a new rule that blocks requests that match the SQL database rule group. Set the web ACL to allow all other traffic that does not match those rules. Attach the web ACL to the ALB in front of the ECS tasks.

D. Create a new AWS WAF web ACL. Create a new empty IP set in AWS WAF. Add a new rule to the web ACL to block requests that originate from IP addresses in the new IP set. Create an AWS Lambda function that scrapes the API logs for IP addresses that send SQL injection attacks, and add those IP addresses to the IP set. Attach the web ACL to the ALB in front of the ECS tasks.

Suggested Answer: *C*

Community vote distribution

C (100%)

dkx Highly Voted 12 months ago

C. Yes, because The SQL database rule group contains rules to block request patterns associated with exploitation of SQL databases, like SQL injection attacks. This can help prevent remote injection of unauthorized queries. Evaluate this rule group for use if your application interfaces with an SQL database.
https://docs.aws.amazon.com/waf/latest/developerguide/aws-managed-rule-groups-use-case.html

A. No, because this does not prevent SQL injection attacks from reaching the ECS API service

B. No, because with Bot Control, you can easily monitor, block, or rate limit bots such as scrapers, scanners, crawlers, status monitors, and search engines.
https://docs.aws.amazon.com/waf/latest/developerguide/waf-bot-control.html

D. No, because because this is a reactive response after a SQL injection attack has occurred for new IP addresses
upvoted 10 times

career360guru Most Recent 7 months, 1 week ago

Selected Answer: C

Option C
upvoted 1 times

NikkyDicky 11 months, 4 weeks ago

Selected Answer: C

C 100%
upvoted 1 times

pupsik 1 year ago

Selected Answer: C

C for sure

upvoted 1 times

Alabi 1 year ago

Selected Answer: C

C for sure

upvoted 1 times

nexus2020 1 year ago

Selected Answer: C

C; the wording is bad. rule is block, and then set the acl to allow everything else that is not matching the block rule?

B: if attacker knows what to attach, coming from a legitment IP, B will not be able to block it, but C can.
D is crazy

upvoted 3 times

Snape 1 year, 1 month ago

Selected Answer: C

Adding new rule for blocking requests which matches SQL database rule group is more 'operationally efficient' than manually scraping API logs and IP based blocking.

upvoted 3 times

ShinLi 1 year, 1 month ago

why not B?

upvoted 1 times

AMEJack 1 year, 1 month ago

Selected Answer: C

Answer is C

upvoted 1 times

Roontha 1 year, 1 month ago

Answer : C
https://docs.aws.amazon.com/waf/latest/developerguide/aws-managed-rule-groups-use-case.html

upvoted 4 times

deegadaze1 1 year, 1 month ago

B- is correct---> AWS WAF Bot Control

upvoted 1 times

Question #184 *Topic 1*

An environmental company is deploying sensors in major cities throughout a country to measure air quality. The sensors connect to AWS IoT Core to ingest timeseries data readings. The company stores the data in Amazon DynamoDB.

For business continuity, the company must have the ability to ingest and store data in two AWS Regions.

Which solution will meet these requirements?

A. Create an Amazon Route 53 alias failover routing policy with values for AWS IoT Core data endpoints in both Regions Migrate data to Amazon Aurora global tables.

B. Create a domain configuration for AWS IoT Core in each Region. Create an Amazon Route 53 latency-based routing policy. Use AWS IoT Core data endpoints in both Regions as values. Migrate the data to Amazon MemoryDB for Redis and configure cross-Region replication.

C. Create a domain configuration for AWS IoT Core in each Region. Create an Amazon Route 53 health check that evaluates domain configuration health. Create a failover routing policy with values for the domain name from the AWS IoT Core domain configurations. Update the DynamoDB table to a global table.

D. Create an Amazon Route 53 latency-based routing policy. Use AWS IoT Core data endpoints in both Regions as values. Configure DynamoDB streams and cross-Region data replication.

Suggested Answer: *C*

Community vote distribution

C (100%)

F_Eldin Highly Voted 2 years, 1 month ago

Selected Answer: C

https://aws.amazon.com/solutions/implementations/disaster-recovery-for-aws-iot/

A, B Wrong. No need to replace DynamoDB with any other DB. DynamoDB Global Table is enough

D- Wrong, Not a use-case for Change Data Capture through Streams

upvoted 9 times

ShenYuying 11 months, 2 weeks ago

The above URL is not available now. You can refer to this URL: https://aws.amazon.com/blogs/iot/how-to-implement-a-disaster-recovery-solution-for-iot-platforms-on-aws/

upvoted 1 times

JosephDZhou Most Recent 1 year, 5 months ago

For C, how failover routing policy have the ability to ingest and store data in two AWS Regions, there is only one active record

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C Business continuity = Failover -> DynamoDB Global DB

upvoted 4 times

NikkyDicky 1 year, 11 months ago

Selected Answer: C

its a C

upvoted 3 times

Maria2023 2 years ago

Selected Answer: C

The only answer which configures DynamoDB properly for multi-region is C

upvoted 2 times

rbm2023 2 years, 1 month ago

Selected Answer: C

Removed B because is replacing Dynamo, unnecessary

upvoted 2 times

andreitugui 2 years, 1 month ago

Selected Answer: C

Answer is C

upvoted 2 times

Roontha 2 years, 1 month ago

Answer: C

upvoted 1 times

Question #185 *Topic 1*

A company uses AWS Organizations for a multi-account setup in the AWS Cloud. The company's finance team has a data processing application that uses AWS Lambda and Amazon DynamoDB. The company's marketing team wants to access the data that is stored in the DynamoDB table.

The DynamoDB table contains confidential data. The marketing team can have access to only specific attributes of data in the DynamoDB table. The finance team and the marketing team have separate AWS accounts.

What should a solutions architect do to provide the marketing team with the appropriate access to the DynamoDB table?

A. Create an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table. Attach the SCP to the OU of the finance team.

B. Create an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes (fine-grained access control). Establish trust with the marketing team's account. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

C. Create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). Attach the policy to the DynamoDB table. In the marketing team's account, create an IAM role that has permissions to access the DynamoDB table in the finance team's account.

D. Create an IAM role in the finance team's account to access the DynamoDB table. Use an IAM permissions boundary to limit the access to the specific attributes. In the marketing team's account, create an IAM role that has permissions to assume the IAM role in the finance team's account.

Suggested Answer: *B*

Community vote distribution

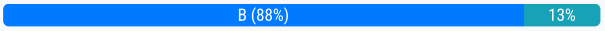

andreitugui Highly Voted 2 years, 1 month ago

Selected Answer: B

Answer is B

upvoted 9 times

pk0619 Most Recent 6 months, 1 week ago

Selected Answer: C

B was right answer until DynamoDB started supporting resource based policies, which makes C right.

upvoted 1 times

liuliangzhou 9 months, 3 weeks ago

Selected Answer: B

I prefer B over C. Attach the policy (specific DynamoDB attributes) to the DynamoDB table. This will result in the finance team's account not being able to fully access DynamoDB.

upvoted 1 times

fartosh 1 year, 1 month ago

Selected Answer: C

I choose C over B.

Both solutions work and are standard approaches for allowing cross-account access. But as compared to S3, option C allows the marketing account to use their usual IAM identities without compromising their permissions. When you assume the role in a different account (option B), you can no longer access resources in your own account.

The resource-based policy for the DynamoDB table supports conditions as well:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/rbac-examples.html#rbac-examples-cross-account

upvoted 2 times

helloworldabc 10 months ago

just B

upvoted 1 times

kgpoj 10 months ago

Dude stop generating garbage info for everyone. I've seen you replying a lot of `just X`. If you have a reason for some choice, then write it down. `just x` sounds so dumb and premature.
upvoted 12 times

sse69 1 year, 1 month ago

Selected Answer: B

Starting march 24', DynamoDB supports resource based policies :
https://aws.amazon.com/about-aws/whats-new/2024/03/amazon-dynamodb-resource-based-policies/
So another way to achieve this would be to create an index for the marketing team, and have the policy restrict their role to that particular index.
On the one hand the new index would incur more costs, on the other hand, having only certain attributes fetched would mean less read units consumed...
upvoted 3 times

yuliaqwerty 1 year, 6 months ago

B https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_dynamodb_attributes.html
upvoted 3 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B as DynamoDB does not support Resource based policies.
upvoted 2 times

LuongTo 7 months ago

Amazon DynamoDB now supports resource-based policies from Mar 20, 2014 https://aws.amazon.com/about-aws/whats-new/2024/03/amazon-dynamodb-resource-based-policies/
upvoted 1 times

erenbiku1 1 year, 6 months ago

Service-linked roles for DynamoDB is not supported
Service roles for DynamoDB is supported
Identity-based policies for DynamoDB is supported
Resource-based policies within DynamoDB is not supported
upvoted 1 times

AMohanty 1 year, 10 months ago

For Cross Account permission we attach Resource Policy with Principal identified as incoming Request Account ARN
+ IAM permissions to query the Finance Account.
C seems more of a resonable answer.
upvoted 1 times

chikorita 1 year, 9 months ago

i dont think C can address the requirement of "he marketing team can have access to only specific attributes of data in the DynamoDB table" hence, B
upvoted 1 times

ggrodskiy 1 year, 11 months ago

Correct C.
upvoted 1 times

Gmail78 1 year, 10 months ago

While resource-based policies can provide granular access control, they are typically used for controlling access within the same AWS account. Cross-account access control is typically achieved using IAM roles with trust relationships. It is B.
upvoted 1 times

AMohanty 1 year, 10 months ago

No, Resource based policies can specify which Principals to give access to Cross Account.
upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

B. DynamoDB fine-grained access using IAM
upvoted 1 times

SkyZeroZx 1 year, 12 months ago

Selected Answer: B

B for sure.
Key word: trust
upvoted 3 times

Maria2023 2 years ago

Selected Answer: B

D would be the perfect choice, since the boundaries are the "new fancy thing" but it's lacking the trust to the marketing account which is a requirement to assume role from one account to another. So it should be B
upvoted 3 times

0c118eb 1 year, 6 months ago

This would not be a good use case for permissions boundaries by itself. Even with permissions boundaries you would still need to implement a solution like B to provide the required permissions.
upvoted 1 times

Alabi 2 years ago

Selected Answer: B

B for sure.
Key word: trust
upvoted 3 times

kfrum4 2 years ago

Selected Answer: B

Answer: B

DynamoDB doesn't support resource based policy
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/using-identity-based-policies.html
upvoted 2 times

ggrodskiy 1 year, 11 months ago

That is not correct. DynamoDB does support resource-based policies for tables and indexes. You can attach a resource-based policy to a DynamoDB table or index to specify who can access that resource and under what conditions. You can also use resource-based policies to grant cross-account access or fine-grained access control for specific DynamoDB attributes. For more information, please refer to this documentation: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/using-identity-based-policies.html
upvoted 1 times

Rajivjain 2 years ago

Selected Answer: C

Resource-based IAM policy
upvoted 1 times

Roontha 2 years, 1 month ago

Answer : B
upvoted 2 times

Question #186 *Topic 1*

A solutions architect is creating an application that stores objects in an Amazon S3 bucket. The solutions architect must deploy the application in two AWS Regions that will be used simultaneously. The objects in the two S3 buckets must remain synchronized with each other.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose three.)

A. Create an S3 Multi-Region Access Point Change the application to refer to the Multi-Region Access Point

B. Configure two-way S3 Cross-Region Replication (CRR) between the two S3 buckets

C. Modify the application to store objects in each S3 bucket

D. Create an S3 Lifecycle rule for each S3 bucket to copy objects from one S3 bucket to the other S3 bucket

E. Enable S3 Versioning for each S3 bucket

F. Configure an event notification for each S3 bucket to invoke an AWS Lambda function to copy objects from one S3 bucket to the other S3 bucket

Suggested Answer: *ABE*

Community vote distribution

ABE (100%)

chathur Highly Voted 1 year ago

Selected Answer: ABE

A - Multi Region Access points are like a proxy. It can dynamically request traffic to the nearest S3 bucket (latency based). [1]

B - Two way replication must be enabled to have data in sync. [1]
E - Versioning must be enabled for Replication. [3]

[1] https://aws.amazon.com/s3/features/multi-region-access-points/
[2] https://aws.amazon.com/about-aws/whats-new/2020/12/amazon-s3-replication-adds-support-two-way-replication/
[3] https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html#two-way-replication-scenario:~:text=Both%20source%20and%20destination%20buckets%20must%20have%20versioning%20enabled.%20For%20more%20information%20about%

upvoted 17 times

SkyZeroZx Highly Voted 12 months ago

Selected Answer: ABE

Cross Region Replication(CRR) requires versioning to be activated due to the way that data is replicated between S3 buckets.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/MultiRegionAccessPointRequestRouting.html

https://stackoverflow.com/questions/60947157/aws-s3-replication-without-versioning#:~:text=The%20automated%20Same%20Region%20Replication,is%20replicated%20between%20S3%20buckets.

upvoted 7 times

career360guru Most Recent 7 months, 1 week ago

Selected Answer: ABE

A, B, E

upvoted 1 times

SK_Tyagi 10 months, 1 week ago

Selected Answer: ABE

Reason as explained by everyone

upvoted 1 times

NikkyDicky 11 months, 4 weeks ago

Selected Answer: ABE

ABE for sure

upvoted 1 times

rbm2023 1 year, 1 month ago

Selected Answer: ABE

I only chosen E because the other options were not making much sense. I guess we need versioning in order to use two-way replication.

upvoted 3 times

Jesuisleon 1 year ago

yes, Cross Region Replication can be implemented only when the versioning of both the buckets is enabled.

upvoted 2 times

Snape 1 year, 1 month ago

Selected Answer: ABE

A. Create an S3 Multi-Region Access Point. - this gives you Single Endpoint for accessing S3 into multiple regions

B. Configure CRR between the two S3 - For automatic replication to diffrent region

E. Enable S3 Versioning on both S3 - Will give you an ability to track and recover from previous versions if needed

C, D and F doesnt meet the criteria from LEAST operation overhead perspective.

upvoted 5 times

F_Eldin 1 year, 1 month ago

Selected Answer: ABE

If the reason for E is not obvious then read this:

https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html

Both source and destination buckets must have versioning enabled.

upvoted 3 times

Bobbyyy 1 year, 1 month ago

Cross Region Replication(CRR) requires versioning to be activated due to the way that data is replicated between S3 buckets.

https://docs.aws.amazon.com/AmazonS3/latest/userguide/MultiRegionAccessPointRequestRouting.html

https://stackoverflow.com/questions/60947157/aws-s3-replication-without-versioning#:~:text=The%20automated%20Same%20Region%20Replication,is%20replicated%20between%20S3%20buckets.

upvoted 1 times

AMEJack 1 year, 1 month ago

Selected Answer: ABE

Answer is A B E

upvoted 1 times

Roontha 1 year, 1 month ago

Answer : A,B,E

upvoted 2 times

Question #187 *Topic 1*

A company has an IoT platform that runs in an on-premises environment. The platform consists of a server that connects to IoT devices by using the MQTT protocol. The platform collects telemetry data from the devices at least once every 5 minutes. The platform also stores device metadata in a MongoDB cluster.

An application that is installed on an on-premises machine runs periodic jobs to aggregate and transform the telemetry and device metadata. The application creates reports that users view by using another web application that runs on the same on-premises machine. The periodic jobs take 120-600 seconds to run. However, the web application is always running.

The company is moving the platform to AWS and must reduce the operational overhead of the stack.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose three.)

A. Use AWS Lambda functions to connect to the IoT devices

B. Configure the IoT devices to publish to AWS IoT Core

C. Write the metadata to a self-managed MongoDB database on an Amazon EC2 instance

D. Write the metadata to Amazon DocumentDB (with MongoDB compatibility)

E. Use AWS Step Functions state machines with AWS Lambda tasks to prepare the reports and to write the reports to Amazon S3. Use Amazon CloudFront with an S3 origin to serve the reports

F. Use an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with Amazon EC2 instances to prepare the reports. Use an ingress controller in the EKS cluster to serve the reports

Suggested Answer: *BDE*

Community vote distribution

BDE (100%)

rbm2023 Highly Voted 1 year, 1 month ago

Selected Answer: BDE

Not A - lambda to connect to IoT is no good

Not C - ec2 instance to run MongoDB

E or F - the job should be short 600 seconds top and serve the reports using Cloud Front - E

upvoted 6 times

career360guru Most Recent 7 months, 1 week ago

Selected Answer: BDE

B, D, E

upvoted 2 times

SK_Tyagi 10 months, 1 week ago

Selected Answer: BDE

F is EKS on EC2 and question is Least Operational overhead

upvoted 3 times

softarts 10 months, 3 weeks ago

E=> how does step function run periodic jobs?

upvoted 1 times

ggrodskiy 11 months, 1 week ago

Correct BDE.

upvoted 1 times

NikkyDicky 11 months, 4 weeks ago

Selected Answer: BDE

BDE for sure

upvoted 2 times

andreitugui 1 year, 1 month ago

Selected Answer: BDE

Answer is B D E

upvoted 1 times

AMEJack 1 year, 1 month ago

Selected Answer: BDE

Support B D E

upvoted 3 times

Roontha 1 year, 1 month ago

Answer : B,D,E

https://aws.amazon.com/step-functions/use-cases/

upvoted 4 times

deegadaze1 1 year, 1 month ago

Correct is ABD

upvoted 1 times

ShinLi 1 year, 1 month ago

why E is wrong?

upvoted 1 times

Question #188 *Topic 1*

A global manufacturing company plans to migrate the majority of its applications to AWS. However, the company is concerned about applications that need to remain within a specific country or in the company's central on-premises data center because of data regulatory requirements or requirements for latency of single-digit milliseconds. The company also is concerned about the applications that it hosts in some of its factory sites, where limited network infrastructure exists.

The company wants a consistent developer experience so that its developers can build applications once and deploy on premises, in the cloud, or in a hybrid architecture. The developers must be able to use the same tools, APIs, and services that are familiar to them.

Which solution will provide a consistent hybrid experience to meet these requirements?

A. Migrate all applications to the closest AWS Region that is compliant. Set up an AWS Direct Connect connection between the central on-premises data center and AWS. Deploy a Direct Connect gateway.

B. Use AWS Snowball Edge Storage Optimized devices for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds. Retain the devices on premises. Deploy AWS Wavelength to host the workloads in the factory sites.

C. Install AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds. Use AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites.

D. Migrate the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds to an AWS Local Zone. Deploy AWS Wavelength to host the workloads in the factory sites.

Suggested Answer: *C*

Community vote distribution

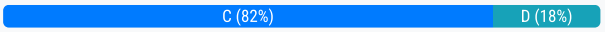

geoakes Highly Voted 1 year, 1 month ago

Selected Answer: C

Key comment: "specific country or in the company's central on-premises data center because of data regulatory requirements or requirements for latency of single-digit milliseconds."

A - No - Region doesn't assure you have in country presence for data soverignty

B - No - Snowball part is correct. However, Wavelength access is only via mobile networks, and not in every country, so this is not possible unless all developers are connecting over the mobile network that will have speed variations

D - No - Local Zones can be fast with a DX connection, but this option like Wavelenght is not in every country

Correct answer is C. 100% of the time you are on premise providing single-digit milliseconds latency as Outposts (rack or server) and Snowball will be in the country for the requirements

upvoted 13 times

pk0619 Most Recent 6 months, 1 week ago

Selected Answer: D

Local zone provides that low latency without having to manage the infrastructure

upvoted 1 times

career360guru 7 months, 1 week ago

Selected Answer: C

Option C

upvoted 1 times

SK_Tyagi 10 months, 1 week ago

Selected Answer: C

Wavelength doesn't makes sense here

upvoted 1 times

NikkyDicky 11 months, 4 weeks ago

Selected Answer: C

C works

upvoted 1 times

pupsik 1 year ago

Selected Answer: C

Wasn't sure abut Snowball Edge compute optimized to run workloads, but it appears to be quite capable option.
Ref: https://docs.aws.amazon.com/snowball/latest/developer-guide/whatisedge.html#edge-related

upvoted 2 times

rbm2023 1 year, 1 month ago

Selected Answer: C

short decision based on brief search
Not B nor D - https://aws.amazon.com/wavelength/
A will not meet the millisecond requirement

upvoted 1 times

Nash101 1 year, 1 month ago

Answer C

Installing AWS Outposts for the applications that have data regulatory requirements or requirements for latency of single-digit milliseconds will provide a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer premises1. AWS Outposts allows customers to run some AWS services locally and connect to a broad range of services available in the local AWS Region1. Using AWS Snowball Edge Compute Optimized devices to host the workloads in the factory sites will provide local compute and storage resources for locations with limited network infrastructure2. AWS Snowball Edge devices can run Amazon EC2 instances and AWS Lambda functions locally and sync data with AWS when network connectivity is available2.

upvoted 3 times

Roontha 1 year, 1 month ago

Answer : C

Reference : https://aws.amazon.com/blogs/compute/aws-local-zones-and-aws-outposts-choosing-the-right-technology-for-your-edge-workload/#:~:text=Unlike%20Outposts%2C%20which%20you%20deploy,using%20for%20an%20AWS%20Region

Local Zones and Outposts can both help you achieve low latency for their latency sensitive workloads. With Direct Connect available in Local Zones, you can achieve low single-digit millisecond latencies, require for applications in online gaming, Media and Entertainment, some SaaS services, AR and VR content delivery etc.

Because Outposts are installed on premises of customers or their data centers, you can achieve under 1 millisecond latencies for workloads that require it.

upvoted 2 times

Masonyeoh 1 year, 1 month ago

Selected Answer: D

Local Zone reduce the latency issue

upvoted 4 times

geoakes 1 year, 1 month ago

Yes, a local zone reduces latency, but local zone are not in every country. The closest thing to an every country option is Snowball and Outpost

upvoted 1 times

Roontha 1 year, 1 month ago

@Masonyeoh, can you review this aws information page on local zones and outposts, confirm your answer again.

https://aws.amazon.com/blogs/compute/aws-local-zones-and-aws-outposts-choosing-the-right-technology-for-your-edge-workload/#:~:text=Unlike%20Outposts%2C%20which%20you%20deploy,using%20for%20an%20AWS%20Region.

upvoted 1 times

ShinLi 1 year, 1 month ago

https://docs.aws.amazon.com/wavelength/latest/developerguide/what-is-wavelength.html

upvoted 1 times

geoakes 1 year, 1 month ago

Wavelength is not present is every country with a datacenter, so B and D options are automatically wrong

upvoted 1 times

Roontha 1 year ago

Answer : C

https://aws.amazon.com/blogs/compute/aws-local-zones-and-aws-outposts-choosing-the-right-technology-for-your-edge-workload/#:~:text=Unlike%20Outposts%2C%20which%20you%20deploy,using%20for%20an%20AWS%20Region.

What is Outposts?
Outposts is a family of fully managed solutions delivering AWS infrastructure and services to virtually any on-premises or edge location for a truly consistent hybrid experience.
upvoted 1 times

Question #189 *Topic 1*

A company is updating an application that customers use to make online orders. The number of attacks on the application by bad actors has increased recently.

The company will host the updated application on an Amazon Elastic Container Service (Amazon ECS) cluster. The company will use Amazon DynamoDB to store application data. A public Application Load Balancer (ALB) will provide end users with access to the application. The company must prevent attacks and ensure business continuity with minimal service interruptions during an ongoing attack.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

A. Create an Amazon CloudFront distribution with the ALB as the origin. Add a custom header and random value on the CloudFront domain. Configure the ALB to conditionally forward traffic if the header and value match.

B. Deploy the application in two AWS Regions. Configure Amazon Route 53 to route to both Regions with equal weight.

C. Configure auto scaling for Amazon ECS tasks Create a DynamoDB Accelerator (DAX) cluster.

D. Configure Amazon ElastiCache to reduce overhead on DynamoDB.

E. Deploy an AWS WAF web ACL that includes an appropriate rule group. Associate the web ACL with the Amazon CloudFront distribution.

Suggested Answer: *AE*

Community vote distribution

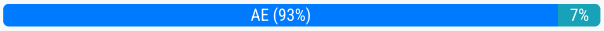

Jackhemo Highly Voted 2 years ago

Selected Answer: AE

From Olabiba.ai:

Option A: By adding a custom header and random value on the CloudFront domain and configuring the ALB to conditionally forward traffic if the header and value match, you can implement a form of request validation. This helps to filter out potentially malicious requests and prevent attacks from reaching the application.
- Option E: Deploying an AWS WAF web ACL that includes an appropriate rule group and associating it with the Amazon CloudFront distribution adds an additional layer of protection. The web ACL can include rules to block common attack patterns and provide protection against various types of attacks, such as SQL injection and cross-site scripting (XSS).

upvoted 6 times

sammyhaj Most Recent 6 months, 3 weeks ago

Selected Answer: BE

broken question

B has business continuity
E must be chosen
A has no business continuity, just recovery or mitigation

upvoted 2 times

43c89f4 1 year, 2 months ago

simple BCD are not at all related to question

upvoted 1 times

Russs99 1 year, 6 months ago

Selected Answer: BE

none of the previous responses really make use of Business continuity as indicated in the scenario. my picks are options B and E. The combination of these two options (E and B) provides both security (via AWS WAF) and high availability (via multi-region deployment) for your application. It helps in preventing attacks and ensuring business continuity with minimal service interruptions during ongoing attacks, making it a cost-effective choice.

upvoted 3 times

kejam 1 year, 5 months ago

Can't use E without A. E depends on A for the CloudFront distribution.

upvoted 7 times

career360guru 1 year, 7 months ago

Selected Answer: AE

A and E

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: AE

its AE

upvoted 2 times

SkyZeroZx 2 years ago

Selected Answer: AE

The only options that helps to protect are A E

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: AE

its a combination of steps, only two of them mention cloud front A and E. it would also be the cheapest option to protect against attacks without having to increase unnecessary performance to the infrastructure which would only cost more money (setup additional region - B , configure auto scaling for ECS and add a DAX - C, configure caching , D).

upvoted 4 times

andreitugui 2 years, 1 month ago

Selected Answer: AE

The only options that helps to protect are A E

upvoted 2 times

Roontha 2 years, 1 month ago

Answer : A E

upvoted 1 times

Question #190 *Topic 1*

A company runs a web application on AWS. The web application delivers static content from an Amazon S3 bucket that is behind an Amazon CloudFront distribution. The application serves dynamic content by using an Application Load Balancer (ALB) that distributes requests to a fleet of Amazon EC2 instances in Auto Scaling groups. The application uses a domain name setup in Amazon Route 53.

Some users reported occasional issues when the users attempted to access the website during peak hours. An operations team found that the ALB sometimes returned HTTP 503 Service Unavailable errors. The company wants to display a custom error message page when these errors occur. The page should be displayed immediately for this error code.

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

A. Set up a Route 53 failover routing policy. Configure a health check to determine the status of the ALB endpoint and to fail over to the failover S3 bucket endpoint.

B. Create a second CloudFront distribution and an S3 static website to host the custom error page. Set up a Route 53 failover routing policy. Use an active-passive configuration between the two distributions.

C. Create a CloudFront origin group that has two origins. Set the ALB endpoint as the primary origin. For the secondary origin, set an S3 bucket that is configured to host a static website Set up origin failover for the CloudFront distribution. Update the S3 static website to incorporate the custom error page.

D. Create a CloudFront function that validates each HTTP response code that the ALB returns. Create an S3 static website in an S3 bucket. Upload the custom error page to the S3 bucket as a failover. Update the function to read the S3 bucket and to serve the error page to the end users.

Suggested Answer: *C*

Community vote distribution

C (60%) D (40%)

pupsik Highly Voted 2 years ago

Selected Answer: C

Origin Groups in CloudFront is what we need here.

upvoted 6 times

Jackhemo Highly Voted 2 years ago

Selected Answer: C

From olabiba.ai:

By using a CloudFront origin group with two origins, you can configure failover between the ALB endpoint and the S3 bucket hosting the static website. This ensures that if the ALB returns HTTP 503 Service Unavailable errors, CloudFront will automatically failover to the S3 bucket and serve the custom error page.

Setting up origin failover for the CloudFront distribution allows for immediate failover to the secondary origin when the primary origin is unavailable. This minimizes the impact of the ALB errors and provides a seamless experience for users by displaying the custom error page.

Updating the S3 static website to incorporate the custom error page ensures that the error page is readily available and can be served to users without any additional processing or delays.

upvoted 5 times

chris_spencer Most Recent 8 months, 3 weeks ago

Selected Answer: C

C because of custom error pages

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/creating-custom-error-pages.html

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: D

A and B are plainly wrong and can be eliminated straight away. The choice therefore is between C and D. The question asks for an immediate display of a custom error page - NOT about permanent failover. Therefore, the correct answer is D.

upvoted 4 times

altonh 5 months, 1 week ago

D is wrong because of this statement: "Update the function to read the S3 bucket and serve the error page to the end users."
CloudFront function cannot do any network access.

upvoted 1 times

fartosh 1 year, 1 month ago

According to https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/high_availability_origin_failover.html
CloudFront always tries to serve the content from the primary origin first.
> CloudFront routes all incoming requests to the primary origin, even when a previous request failed over to the secondary origin. CloudFront only sends requests to the secondary origin after a request to the primary origin fails.

Therefore option C is still valid as it does not leave CloudFront in "permanent failover".

upvoted 1 times

chelbsik 1 year, 4 months ago

Selected Answer: D

I go for D: it contains all steps to setup the requested solution, and CloudFront function suits here
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cloudfront-functions.html
"URL redirects or rewrites – You can redirect viewers to other pages based on information in the request, or rewrite all requests from one path to another".

upvoted 1 times

AimarLeo 1 year, 4 months ago

Selected Answer: D

'The company wants to display a custom error message page when these errors occur. The page should be displayed immediately for this error code.'
The purpose of the question obviously is to return that error page not really a FAILOVER mechanism --> Leaves D as an asnwer

upvoted 3 times

carpa_jo 1 year, 6 months ago

For people are asking why C is better than A:
The approach of A is more suited for scenarios where there is a complete failure of the primary endpoint rather than intermittent errors. The health checks may not register a failure if the 502 errors are sporadic and the system is generally operational, thus the failover might not be triggered.
With the approach of C CloudFront will always automatically switch to the secondary origin when the primary origin returns specific HTTP status code failure responses.

upvoted 3 times

Niko13 1 year, 6 months ago

Selected Answer: C

Least Operational Overhead is C

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: C

Least Operational Overhead is C

upvoted 2 times

KCjoe 1 year, 8 months ago

I know C is good, but why not A, seems to me A is much easier.

upvoted 1 times

SuperDuperPooperScooper 1 year, 7 months ago

Route 53 failover will not be as immediate as C. Cloudfront will immediately seerve up the error page if the request to the primary origin fails, so there is no delay between the primary origin health being degraded and the failover page being served.

upvoted 2 times

bur4an 1 year, 9 months ago

Repeat question?

upvoted 1 times

kjcncjek 1 year, 10 months ago

why not A?

upvoted 3 times

hamimelon 8 months, 3 weeks ago

Route 53 fail over to S3? How can Route 53 display the image?

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: C

it's a C

upvoted 2 times

rbm2023 2 years, 1 month ago

Almost went for D but this would take too much operational overhead.

upvoted 2 times

rbm2023 2 years, 1 month ago

Option C

upvoted 1 times

andreitugui 2 years, 1 month ago

Selected Answer: C

Answer is C, you can use origin groups and configure error response pages in Cloud Front based on different request response codes (503, 404, 403 etc)

upvoted 3 times

Roontha 2 years, 1 month ago

Answer : C

https://repost.aws/knowledge-center/cloudfront-distribution-serve-content

upvoted 3 times

Question #191 *Topic 1*

A company is planning to migrate an application to AWS. The application runs as a Docker container and uses an NFS version 4 file share.

A solutions architect must design a secure and scalable containerized solution that does not require provisioning or management of the underlying infrastructure.

Which solution will meet these requirements?

A. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type. Use Amazon Elastic File System (Amazon EFS) for shared storage. Reference the EFS file system ID, container mount point, and EFS authorization IAM role in the ECS task definition.

B. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type. Use Amazon FSx for Lustre for shared storage. Reference the FSx for Lustre file system ID, container mount point, and FSx for Lustre authorization IAM role in the ECS task definition.

C. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type and auto scaling turned on. Use Amazon Elastic File System (Amazon EFS) for shared storage. Mount the EFS file system on the ECS container instances. Add the EFS authorization IAM role to the EC2 instance profile.

D. Deploy the application containers by using Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type and auto scaling turned on. Use Amazon Elastic Block Store (Amazon EBS) volumes with Multi-Attach enabled for shared storage. Attach the EBS volumes to ECS container instances. Add the EBS authorization IAM role to an EC2 instance profile.

Suggested Answer: *A*

Community vote distribution

A (100%)

chris_spencer 8 months, 3 weeks ago

It's very easy... you read docker => ECS. NFS => EFS, no underlaying infrastructure => Fargate

upvoted 1 times

saggy4 1 year, 4 months ago

Selected Answer: A

C and D: Both these options have hassles of EC2 management

Between A and B: Mounting FSx for Lustre on an AWS Fargate launch type isn't supported.

Hence the correct option is A

upvoted 3 times

Niko13 1 year, 6 months ago

Selected Answer: A

ECS, EFS - answer A

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: A

Option A -

EFS = NFS 4

Fargate = No mgmt or provisioning overheads for servers

upvoted 3 times

Christina666 1 year, 11 months ago

Selected Answer: A

Amazon EFS is a managed NAS filer for EC2 instances based on Network File System (NFS) version 4.

upvoted 4 times

NikkyDicky 1 year, 11 months ago

Selected Answer: A

A for sure

upvoted 1 times

SkyZeroZx 1 year, 12 months ago

Selected Answer: A

A is correct

B Fsx For Lustre is POSIX Compilance not is correct in this question

C and D usage EC2 more overhead administrative is incorrect

upvoted 2 times

Gishpi 1 year, 11 months ago

EFS is POSIX Compliant too. A is correct, because EFS file systems can be accessed by Amazon EC2 Linux instances, Amazon ECS, Amazon EKS, AWS Fargate, and AWS Lambda functions via a file system interface such as NFS protocol.

upvoted 2 times

Maria2023 2 years ago

Selected Answer: A

https://aws.amazon.com/fsx/when-to-choose-fsx/

upvoted 2 times

rbm2023 2 years, 1 month ago

Selected Answer: A

Must be fargate due to the "not require provisioning or management of the underlying infra"

A or B , tie breaker using EFS and not FSx

Hence option A.

upvoted 1 times

andreitugui 2 years, 1 month ago

Selected Answer: A

The correct answer is A, fargate(no infra management) & efs for NFSv4

upvoted 2 times

deegadaze1 2 years, 1 month ago

A is correct due to -- NFS version 4.

upvoted 3 times

Roontha 2 years, 1 month ago

Answer : A

https://aws.amazon.com/about-aws/whats-new/2017/03/amazon-elastic-file-system-amazon-efs-now-supports-nfsv4-lock-upgrading-and-downgrading/

upvoted 1 times

Question #192 *Topic 1*

A company is running an application in the AWS Cloud. The core business logic is running on a set of Amazon EC2 instances in an Auto Scaling group. An Application Load Balancer (ALB) distributes traffic to the EC2 instances. Amazon Route 53 record api.example.com is pointing to the ALB.

The company's development team makes major updates to the business logic. The company has a rule that when changes are deployed, only 10% of customers can receive the new logic during a testing window. A customer must use the same version of the business logic during the testing window.

How should the company deploy the updates to meet these requirements?

A. Create a second ALB, and deploy the new logic to a set of EC2 instances in a new Auto Scaling group. Configure the ALB to distribute traffic to the EC2 instances. Update the Route 53 record to use weighted routing, and point the record to both of the ALBs.

B. Create a second target group that is referenced by the ALDeploy the new logic to EC2 instances in this new target group. Update the ALB listener rule to use weighted target groups. Configure ALB target group stickiness.

C. Create a new launch configuration for the Auto Scaling group. Specify the launch configuration to use the AutoScalingRollingUpdate policy, and set the MaxBatchSize option to 10. Replace the launch configuration on the Auto Scaling group. Deploy the changes.

D. Create a second Auto Scaling group that is referenced by the ALB. Deploy the new logic on a set of EC2 instances in this new Auto Scaling group. Change the ALB routing algorithm to least outstanding requests (LOR). Configure ALB session stickiness.

Suggested Answer: *B*

Community vote distribution

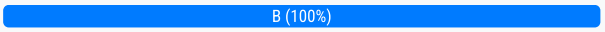

career360guru 7 months, 1 week ago

Selected Answer: B

B is better option considering the fact that a customer should get same business logic during testing window. This means we need session stickiness that only option B can provide.

upvoted 4 times

Pupu86 7 months, 1 week ago

Selected Answer: B

This is canary deployment not blue/green

upvoted 3 times

joleneinthebackyard 8 months ago

Selected Answer: B

I was struggled between A and B because I overlooked this line "A customer must use the same version of the business logic during the testing window."

So we need session stickiness in place, then B is the obvious choice.

upvoted 1 times

aviathor 10 months, 2 weeks ago

The problem I have with B is that is does not mention stickiness. The problem I have with A is that the stickiness will work only as long as the DNS entry does not time out...

upvoted 1 times

aviathor 10 months, 2 weeks ago

Oops. It does mention stickiness...

upvoted 1 times

ggrodskiy 11 months, 1 week ago

Correct B.

upvoted 1 times

NikkyDicky 11 months, 4 weeks ago

Selected Answer: B

B better

upvoted 1 times

SkyZeroZx 1 year ago

Selected Answer: B

B) Classic usage of Blue/Green deployment

A is good option but not have a stickness with Route 53 more apropiate is ALB with stickness

upvoted 2 times

Maria2023 1 year ago

Selected Answer: B

https://docs.aws.amazon.com/prescriptive-guidance/latest/load-balancer-stickiness/target-group-stickiness.html

upvoted 1 times

rbm2023 1 year, 1 month ago

Selected Answer: B

Agree with B

blue green deployment, using target group

upvoted 4 times

rbm2023 1 year, 1 month ago

https://aws.amazon.com/blogs/aws/new-application-load-balancer-simplifies-deployment-with-weighted-target-groups/

upvoted 3 times

F_Eldin 1 year, 1 month ago

Selected Answer: B

https://aws.amazon.com/blogs/aws/new-application-load-balancer-simplifies-deployment-with-weighted-target-groups/

upvoted 4 times

Roontha 1 year, 1 month ago

Answer : B

https://medium.com/capital-one-tech/deploying-with-confidence-strategies-for-canary-deployments-on-aws-7cab3798823e

upvoted 2 times

Question #193 *Topic 1*

A large education company recently introduced Amazon Workspaces to provide access to internal applications across multiple universities. The company is storing user profiles on an Amazon FSx for Windows File Server file system. The file system is configured with a DNS alias and is connected to a self-managed Active Directory. As more users begin to use the Workspaces, login time increases to unacceptable levels.

An investigation reveals a degradation in performance of the file system. The company created the file system on HDD storage with a throughput of 16 MBps. A solutions architect must improve the performance of the file system during a defined maintenance window.

What should the solutions architect do to meet these requirements with the LEAST administrative effort?

A. Use AWS Backup to create a point-in-time backup of the file system. Restore the backup to a new FSx for Windows File Server file system. Select SSD as the storage type. Select 32 MBps as the throughput capacity. When the backup and restore process is completed, adjust the DNS alias accordingly. Delete the original file system.

B. Disconnect users from the file system. In the Amazon FSx console, update the throughput capacity to 32 MBps. Update the storage type to SSD. Reconnect users to the file system.

C. Deploy an AWS DataSync agent onto a new Amazon EC2 instance. Create a task. Configure the existing file system as the source location. Configure a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput as the target location. Schedule the task. When the task is completed, adjust the DNS alias accordingly. Delete the original file system.

D. Enable shadow copies on the existing file system by using a Windows PowerShell command. Schedule the shadow copy job to create a point-in-time backup of the file system. Choose to restore previous versions. Create a new FSx for Windows File Server file system with SSD storage and 32 MBps of throughput. When the copy job is completed, adjust the DNS alias. Delete the original file system.

Suggested Answer: *A*

Community vote distribution

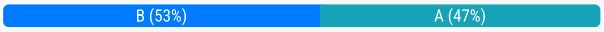

F_Eldin Highly Voted 2 years, 1 month ago

Selected Answer: A

B is wrong :
https://aws.amazon.com/fsx/windows/faqs/#:~:text=A%3A%20While%20you%20cannot%20change,with%20a%20different%20storage%20type.
I can modify the capacity, but not the type.

upvoted 17 times

Sab 1 year, 8 months ago

Storage type can be modified
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/managing-storage-type.html

upvoted 8 times

AK2020 1 year, 8 months ago

You can change your file system storage type from HDD to SSD using the Amazon FSx console or Amazon FSx API. You can't change your file system storage type from SSD to HDD. So A is correct as we can do this during the downtime

upvoted 3 times

AK2020 1 year, 8 months ago

So B is correct. my apologies

upvoted 2 times

Andres123456 Highly Voted 1 year, 7 months ago

Selected Answer: B

Storage type can be modified
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/managing-storage-type.html

upvoted 9 times

0dc6cac Most Recent 2 weeks ago

Selected Answer: B

Tough question, in 10/10 cases I would pick A over B, it's objectively the more appropriate method. However, B technically requires a step or two less than A to perform. It also depends on how much administrative load will be added during the downtime (it can be long, because AWS sometimes

takes forever to stop and start stuff).

So if we assume no time constraint, and no issues with downtime, B is correct. In reality, it's always going to be A.

upvoted 1 times

Kaps443 2 weeks, 5 days ago

Selected Answer: A

B is Incorrect: You cannot change the storage type from HDD to SSD after creation of the FSx file system.

upvoted 1 times

820b83f 4 months, 2 weeks ago

Selected Answer: A

My reasons for its A:

1. FSx does not support live storage type changes from HDD to SSD. You must create a new file system.

upvoted 1 times

bhanus 6 months ago

Selected Answer: B

https://docs.aws.amazon.com/fsx/latest/WindowsGuide/updating-storage-type.html

HDD can be changed to SSd

upvoted 1 times

SIJUTHOMASP 6 months, 1 week ago

Selected Answer: B

Comparing between A and B, the trade-off decision would be on the key requirement for 'least administrative efforts'. Which seems to be lesser in B, because it's single step but on the other hand on A, there are multiple steps of backup, creating new FSx etc should be of more admin efforts. Hence B.

upvoted 1 times

pk0619 6 months, 1 week ago

Selected Answer: B

You can update both throughput capacity as well as storage capacity of an existing filesystem

upvoted 1 times

LuongTo 7 months ago

Selected Answer: B

SSD to HDD is impossible, but HDD to SSD is okay => B is feasible.

B is less effort since B just disconnects users from the file system for a while, and then updates the FSx. While A needs a new FSx, backup, restore, clean up then switch, more steps to do than A

upvoted 2 times

FZA24 7 months, 2 weeks ago

Selected Answer: A

Let consider that B is correct (updating storage type is possible).

Between A and B, A needs the LEAST administrative effort.

A is seamless for users. However, B requires to disconnect users and thus service interruption and administrative effort to manage that!

upvoted 1 times

Sin_Dan 8 months, 1 week ago

Selected Answer: A

I pity those who are selecting B.

updating the storage type (from HDD to SSD) is not supported for an existing FSx for Windows File Server file system. You would need to create a new file system to change the storage type. Therefore, this solution is not feasible.

upvoted 1 times

Zinnia_Wang 6 months, 1 week ago

https://docs.aws.amazon.com/fsx/latest/WindowsGuide/updating-storage-type.html

upvoted 1 times

JoeTromundo 8 months, 3 weeks ago

Selected Answer: B

"You CAN CHANGE your file system storage type from HDD to SSD using the AWS Management Console and AWS CLI."
"You CANNOT CHANGE your file system storage type from SSD to HDD."
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/managing-storage-configuration.html#managing-storage-type
upvoted 1 times

Syre 10 months, 1 week ago

Selected Answer: A

Option B would be incorrect because it mentions updating the throughput and storage type directly in the FSx console, which is not supported for an existing FSx for Windows File Server.
upvoted 1 times

helloworldabc 10 months ago

just B
upvoted 1 times

dragongoseki 1 year ago

Selected Answer: B

B is right answer.
upvoted 3 times

Helpnosense 1 year ago

Selected Answer: B

Since hdd to ssd type is doable. B is better answer.
upvoted 1 times

Bobshaw 1 year, 1 month ago

Selected Answer: A

AWS Backup to create a point-in-time backup of the existing file system, restoring the backup to a new FSx for Windows File Server file system with SSD storage and higher throughput capacity, adjusting the DNS alias, and deleting the original file system provides the most efficient and least administratively intensive solution to improve the performance of the file system during a defined maintenance window
upvoted 2 times

seetpt 1 year, 1 month ago

Selected Answer: B

B is correct
upvoted 1 times

Question #194 *Topic 1*

A company hosts an application on AWS. The application reads and writes objects that are stored in a single Amazon S3 bucket. The company must modify the application to deploy the application in two AWS Regions.

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

A. Set up an Amazon CloudFront distribution with the S3 bucket as an origin. Deploy the application to a second Region Modify the application to use the CloudFront distribution. Use AWS Global Accelerator to access the data in the S3 bucket.

B. Create a new S3 bucket in a second Region. Set up bidirectional S3 Cross-Region Replication (CRR) between the original S3 bucket and the new S3 bucket. Configure an S3 Multi-Region Access Point that uses both S3 buckets. Deploy a modified application to both Regions.

C. Create a new S3 bucket in a second Region Deploy the application in the second Region. Configure the application to use the new S3 bucket. Set up S3 Cross-Region Replication (CRR) from the original S3 bucket to the new S3 bucket.

D. Set up an S3 gateway endpoint with the S3 bucket as an origin. Deploy the application to a second Region. Modify the application to use the new S3 gateway endpoint. Use S3 Intelligent-Tiering on the S3 bucket.

Suggested Answer: *B*

Community vote distribution

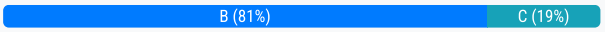

a54b16f 1 year, 3 months ago

Selected Answer: B

C is missing "bidirectional S3 Cross-Region Replication"

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: B

B is always a better option. C is possible but less preferred.

Irrespective of B or C application will need modification to deploy in 2nd region as Bucket URL has to be change in application.

upvoted 3 times

Russs99 1 year, 8 months ago

Selected Answer: C

An S3 Multi-Region Access Point is a global endpoint that provides access to data in one or more S3 buckets. To create an S3 Multi-Region Access Point, you must specify a set of S3 buckets that you want to include in the Multi-Region Access Point. You must also configure routing rules to determine which requests are routed to which S3 buckets.

Once you have created an S3 Multi-Region Access Point, you must modify your application to use the Multi-Region Access Point endpoint instead of the S3 bucket endpoints. This requires changes to your application code and configuration.

Option C does not require the creation of an S3 Multi-Region Access Point. Instead, you can simply deploy the application in two Regions and configure the application to use the S3 bucket endpoints in each Region. This is a simpler and more straightforward approach, which reduces operational overhead.

upvoted 3 times

helloworldabc 10 months ago

just B

upvoted 1 times

carpa_jo 1 year, 6 months ago

Option C includes "Set up S3 Cross-Region Replication (CRR) from the original S3 bucket to the new S3 bucket". By that the application in the new region will have access to the files from the "old" and the new region, and the application running in the "old" region only has access to the data of the "old" region, as no bidirectional CRR is being set up. That doesn't make a lot of sense. Option B contains bidirectional CRR which keeps both buckets in sync.

upvoted 3 times

MasterP007 1 year, 10 months ago

Selected Answer: B

Option B creates a new S3 bucket in a second Region and sets up bidirectional S3 Cross-Region Replication (CRR) between the original S3 bucket and the new S3 bucket. S3 CRR is a feature that enables automatic, asynchronous copying of objects across S3 buckets in different AWS Regions. You can use S3 CRR to keep your data synchronized across Regions for lower latency, compliance, security, disaster recovery, and regional efficiency.

upvoted 4 times

azizmo 1 year, 11 months ago

Selected Answer: B

The answer is B

upvoted 1 times

nicecurls 1 year, 11 months ago

Selected Answer: B

it's a B

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

its a B

upvoted 2 times

NikkyDicky 1 year, 11 months ago

the "stored in a single Amazon S3 bucket" comment is confusing though. have to assume new versionn will have buckets in each region

upvoted 3 times

phattran 1 year, 11 months ago

Selected Answer: B

S3 CRR prefer S3 Multi-Region Access Point

upvoted 3 times

YodaMaster 1 year, 11 months ago

B sounds right for deploying in 2 different regions though.

upvoted 1 times

YodaMaster 1 year, 11 months ago

this question seems incomplete?

upvoted 1 times

Masonyeoh 2 years, 1 month ago

B, enable the S3 sync

upvoted 3 times

Roontha 2 years, 1 month ago

Answer : B

https://aws.amazon.com/s3/features/multi-region-access-points/

upvoted 2 times

Question #195 *Topic 1*

An online gaming company needs to rehost its gaming platform on AWS. The company's gaming application requires high performance computing (HPC) processing and has a leaderboard that changes frequently. An Ubuntu instance that is optimized for compute generation hosts a Node.js application for game display. Game state is tracked in an on-premises Redis instance.

The company needs a migration strategy that optimizes application performance.

Which solution will meet these requirements?

A. Create an Auto Scaling group of m5.large Amazon EC2 Spot Instances behind an Application Load Balancer. Use an Amazon ElastICache for Redis cluster to maintain the leaderboard.

B. Create an Auto Scaling group of c5.large Amazon EC2 Spot Instances behind an Application Load Balancer. Use an Amazon OpenSearch Service cluster to maintain the leaderboard.

C. Create an Auto Scaling group of c5.large Amazon EC2 On-Demand Instances behind an Application Load Balancer. Use an Amazon ElastiCache for Redis cluster to maintain the leaderboard.

D. Create an Auto Scaling group of m5.large Amazon EC2 On-Demand Instances behind an Application Load Balancer. Use an Amazon DynamoDB table to maintain the leaderboard.

Suggested Answer: *C*

Community vote distribution

C (100%)

Roontha Highly Voted 1 year, 7 months ago

Answer : C

https://aws.amazon.com/blogs/database/building-a-real-time-gaming-leaderboard-with-amazon-elasticache-for-redis/

upvoted 12 times

rbm2023 Highly Voted 1 year, 7 months ago

Selected Answer: C

Elastic Cache for Redis, C or D.

Both are on demand, we cant use spot

Tie breaker is the instance type c5.

upvoted 5 times

Win007 Most Recent 6 months, 3 weeks ago

D is the write answer

upvoted 1 times

voccer 11 months, 2 weeks ago

Answer: C

B/c: not use spot instance

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: C

Option C

upvoted 1 times

dkcloudguru 1 year, 3 months ago

Agree with option C

upvoted 1 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: C

Agree with C.

upvoted 1 times

ggrodskiy 1 year, 5 months ago

Correct C.

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

C for sure

upvoted 1 times

YodaMaster 1 year, 5 months ago

Selected Answer: C

C is the way

upvoted 1 times

Alabi 1 year, 6 months ago

Selected Answer: C

C for sure

upvoted 1 times

F_Eldin 1 year, 7 months ago

Selected Answer: C

A, B : Wrong. Spot instances. B: OpeSearch instead of Redis

D: Wrong, DynamoDB instead of Redis

upvoted 2 times

andreitugui 1 year, 7 months ago

Selected Answer: C

The answer is C as compute optimized instance is required c5, and ElastiCache is the for Redis.

upvoted 2 times

Masonyeoh 1 year, 7 months ago

Agree with C

upvoted 2 times

Question #196 *Topic 1*

A solutions architect is designing an application to accept timesheet entries from employees on their mobile devices. Timesheets will be submitted weekly, with most of the submissions occurring on Friday. The data must be stored in a format that allows payroll administrators to run monthly reports. The infrastructure must be highly available and scale to match the rate of incoming data and reporting requests.

Which combination of steps meets these requirements while minimizing operational overhead? (Choose two.)

A. Deploy the application to Amazon EC2 On-Demand Instances with load balancing across multiple Availability Zones. Use scheduled Amazon EC2 Auto Scaling to add capacity before the high volume of submissions on Fridays.

B. Deploy the application in a container using Amazon Elastic Container Service (Amazon ECS) with load balancing across multiple Availability Zones. Use scheduled Service
Auto Scaling to add capacity before the high volume of submissions on Fridays.

C. Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFront. Deploy the application backend using Amazon API Gateway with an AWS Lambda proxy integration.

D. Store the timesheet submission data in Amazon Redshift. Use Amazon QuickSight to generate the reports using Amazon Redshift as the data source.

E. Store the timesheet submission data in Amazon S3. Use Amazon Athena and Amazon QuickSight to generate the reports using Amazon S3 as the data source.

Suggested Answer: *BE*

Community vote distribution

emiliocb4 Highly Voted 2 years ago

Selected Answer: BE

i'm going with BE.
A not correct with EC2 instances to mantain.
C is not correct because we cannot host webapplication on S3 (only static contents)
D too much effort for Redshift
upvoted 18 times

altonh 5 months, 1 week ago

C implies that the front end is static while the back end is dynamic.
upvoted 1 times

Gmail78 1 year, 10 months ago

It looks like BE are the best options. While deploying the frontend to S3 and using API Gateway with Lambda for the backend is a good architectural approach, it might not directly address the requirement for load scaling and scheduling.
upvoted 1 times

YodaMaster Highly Voted 1 year, 11 months ago

Selected Answer: CE

A. EC2 on-demand instances don't make sense to accept timesheet entries
B. ECS can be done but they want to minimise operational overhead where option C sounds better/simple
C. Sounds simple enough to use s3. I choose this.
D. I already chose s3 so this doesn't apply + redshift seems overkill
E.This goes with Option C
So answer C and E
upvoted 11 times

Kaps443 Most Recent 2 weeks, 5 days ago

Selected Answer: CE

C is best for frontend/API
E is best for storage + reports
upvoted 1 times

JaffaDaffa 5 months, 4 weeks ago

Selected Answer: CE

C is better option than B bcz it is managed/serverless than ECS (without mentioning Fargate)

upvoted 1 times

JaffaDaffa 5 months, 4 weeks ago

Selected Answer: CE

B is not right option as ECS management overhead unless specified with Fargate Launch type.

upvoted 1 times

deepakR20 6 months, 3 weeks ago

Selected Answer: BE

Key parameter is " with most of the submissions occurring on Friday." hence BE is the right answer

upvoted 1 times

LuongTo 7 months ago

Selected Answer: BE

E is apparently.

I would go for B rather than C. Even serverless lambda is best suited for minimizing operational overhead; however the point is "mobile devices"; the mobile application is the frontend => "Deploy the application front end to an Amazon S3 bucket served by Amazon CloudFront" from C does not make sense.

upvoted 1 times

youonebe 7 months ago

Selected Answer: CE

CE-CE-CE

upvoted 1 times

zersa 7 months, 2 weeks ago

Selected Answer: BE

i'm going with BE.

upvoted 1 times

Halliphax 7 months, 2 weeks ago

Selected Answer: BE

High availability = multiple availability zones (clue in the answer, B)

Cannot be C as it's a web application so cannot be on S3. Lambda not suitable either because Lambdas only run in a single chosen region of deployment.

upvoted 1 times

LuongTo 7 months ago

C feasible. Frontend (html, js) will be in S3 with Cloudfront, lambda for backend. lambda is the best approach for "minimizing operational overhead". The "requirement" here is about high-availability not DR which requires multi-region

upvoted 2 times

0b43291 7 months, 3 weeks ago

Selected Answer: CE

Option A (EC2 On-Demand Instances with Auto Scaling) requires managing and scaling EC2 instances, which adds operational overhead compared to a serverless approach.

Option B (Amazon ECS with Service Auto Scaling) also requires managing and scaling container instances, which adds operational overhead compared to a serverless approach.

Option D (Amazon Redshift) is a data warehousing solution better suited for large-scale analytics workloads, which may be overkill for the given requirements and introduce unnecessary complexity and cost.

By choosing the combination of options C and E, the solutions architect can implement a highly available, scalable, and cost-effective solution with minimal operational overhead, leveraging the benefits of serverless computing, object storage, and managed analytics services.

upvoted 1 times

sashenka 8 months, 1 week ago

Selected Answer: CE

Options A and B involve managing EC2 instances or containers, which would require more operational effort than a fully serverless solution with C and E.

upvoted 1 times

AWSum1 8 months, 3 weeks ago

Selected Answer: BE

It's B & E, the question says that timeshares will need ro be submitted. Therefore making it dynamic.

Selecting option C to use S3 to host webapp can't work because S3 can host static sites.

upvoted 1 times

Syre 10 months, 1 week ago

Selected Answer: BE

C is Incorrect

upvoted 1 times

kpcert 1 year ago

Selected Answer: BE

Voting for BE

upvoted 2 times

trungtd 1 year ago

Selected Answer: BE

C is not correct. We need Lambda, not Lambda Proxy. BTW, APIGW + Lambda for the unknown load, in this case we knew that high load on Friday

upvoted 2 times

red_panda 1 year, 1 month ago

Selected Answer: BE

Answer is B and E.

We already know which on friday a lot of people will submit timesheet, so scheduled autoscaling is perfect.

Finally S3 + Athena are enough. No need of Redshift database to run report.

upvoted 4 times

Question #197 *Topic 1*

A company is storing sensitive data in an Amazon S3 bucket. The company must log all activities for objects in the S3 bucket and must keep the logs for 5 years. The company's security team also must receive an email notification every time there is an attempt to delete data in the S3 bucket.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose three.)

A. Configure AWS CloudTrail to log S3 data events.

B. Configure S3 server access logging for the S3 bucket.

C. Configure Amazon S3 to send object deletion events to Amazon Simple Email Service (Amazon SES).

D. Configure Amazon S3 to send object deletion events to an Amazon EventBridge event bus that publishes to an Amazon Simple Notification Service (Amazon SNS) topic.

E. Configure Amazon S3 to send the logs to Amazon Timestream with data storage tiering.

F. Configure a new S3 bucket to store the logs with an S3 Lifecycle policy.

Suggested Answer: *ADF*

Community vote distribution

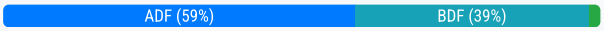

cmoreira Highly Voted 1 year, 9 months ago

Selected Answer: ADF

ADF

A or B work, but docs recomment cloud trail:

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

upvoted 12 times

kaby1987 Highly Voted 1 year, 6 months ago

Selected Answer: ADF

ADF are correct choices.

upvoted 7 times

altonh Most Recent 4 months ago

Selected Answer: BDF

The requirement is for the most cost-effective.

A - You will pay for the data event delivered to S3

C - No integration to SES

D - Paying more for Timestream

upvoted 1 times

altonh 4 months ago

The requirement is for the most cost-effective.

A - You will pay for the data event delivered to S3

C - No integration to SES

E - Paying more for Timestream

upvoted 1 times

820b83f 4 months, 2 weeks ago

Selected Answer: ADF

The conflict is between AWS Cloudtrail (A) and S3 Server Access Log (B).

B is incorrect because S3 Server access logs track requests at the bucket level, not object-level operations (e.g., deletions).

A is correct because CloudTrail is required for detailed tracking including object level..

upvoted 1 times

altonh 5 months, 1 week ago

Selected Answer: BDF

BDF flows well. ADF, however, does not provide details on how you will store the logs in the new S3 bucket.
upvoted 1 times

HSong 9 months, 3 weeks ago
"We recommend that you use CloudTrail for logging bucket-level and object-level actions for your Amazon S3 resources."
upvoted 3 times

dragongoseki 1 year ago
Selected Answer: ADE
ADFis right answer.
upvoted 1 times

Helpnosense 1 year ago
Selected Answer: BDF
Both A and B can log s3 activities. Difference is A real time log but cost more. B log has delay but cheaper. The requirement is the most cost-effective so choose B to meet this requirement.
upvoted 4 times

seetpt 1 year, 1 month ago
Selected Answer: ADF
ADF logs everything, BDF doesnt.
upvoted 3 times

titi_r 1 year, 2 months ago
Selected Answer: BDF
BDF meet the requirements.
upvoted 3 times

liquen14 1 year, 3 months ago
Selected Answer: ADF
Probably B is cheaper but A is safer and more accurate and remember the "The company must log ALL activities for objects"

According to this https://docs.aws.amazon.com/AmazonS3/latest/userguide/ServerLogs.html#LogDeliveryBestEffort

"The log record for a particular request might be delivered long after the request was actually processed, or it might not be delivered at all. "

so for me is A not B
upvoted 5 times

Russs99 1 year, 4 months ago
Selected Answer: BDF
Given the requirement to log all activities for objects in an S3 bucket and keep logs for 5 years, combined with a focus on cost-effectiveness, S3 server access logging (Option B) would indeed be a cheaper solution for capturing basic access logs. However, for advanced auditing and compliance requirements where detailed API call tracking is needed, CloudTrail's data event logging provides valuable insights that S3 access logs do not.
upvoted 4 times

ninomfr64 1 year, 5 months ago
Selected Answer: BDF
B is cheaper than A
AWS CloudTrail (A) - Management events (first delivery) are free; data events incur a fee, in addition to storage of logs
S3 Server Logs (B) - No other cost in addition to storage of logs

https://docs.aws.amazon.com/AmazonS3/latest/userguide/logging-with-S3.html#:~:text=S3%20Server%20Logs-,Price,-Management%20events%20(first
upvoted 1 times

gagol14 1 year, 5 months ago
Selected Answer: ADF
For capturing object-level events, such as object deletions, you would typically use Amazon S3 Event Notifications or enable AWS CloudTrail data events for S3.
upvoted 4 times

Jane1234YIP 1 year, 5 months ago

S3 server access logging does not capture object-level events like object deletions. so I will go ADF.

upvoted 3 times

cox1960 1 year, 5 months ago

wrong. check "operation" in https://docs.aws.amazon.com/AmazonS3/latest/userguide/LogFormat.html

BDF

upvoted 5 times

adelynllllllllll 1 year, 6 months ago

BDF

Because it asked for cost-effective.

upvoted 1 times

mosalahs 1 year, 6 months ago

Selected Answer: BDF

B is better than A because S3 server logs -- > Cost efficient and get more log information (Lifecycle, Authentication info)

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

upvoted 2 times

Question #198 *Topic 1*

A company is building a hybrid environment that includes servers in an on-premises data center and in the AWS Cloud. The company has deployed Amazon EC2 instances in three VPCs. Each VPC is in a different AWS Region. The company has established an AWS Direct. Connect connection to the data center from the Region that is closest to the data center.

The company needs the servers in the on-premises data center to have access to the EC2 instances in all three VPCs. The servers in the on-premises data center also must have access to AWS public services.

Which combination of steps will meet these requirements with the LEAST cost? (Choose two.)

A. Create a Direct Connect gateway in the Region that is closest to the data center. Attach the Direct Connect connection to the Direct Connect gateway. Use the Direct Connect gateway to connect the VPCs in the other two Regions.

B. Set up additional Direct Connect connections from the on-premises data center to the other two Regions.

C. Create a private VIF. Establish an AWS Site-to-Site VPN connection over the private VIF to the VPCs in the other two Regions.

D. Create a public VIF. Establish an AWS Site-to-Site VPN connection over the public VIF to the VPCs in the other two Regions.

E. Use VPC peering to establish a connection between the VPCs across the Regions Create a private VIF with the existing Direct Connect connection to connect to the peered VPCs.

Suggested Answer: *AD*

Community vote distribution

AD (100%)

cmoreira Highly Voted 1 year, 9 months ago

Selected Answer: AD

There is no correct answer. NONE.
A.Direct Connect gateway are global. You dont create them in a "region"
B. Not needed, since you have DX-GW.
C. Cant establish site-to-site VPN over private VIF. You do it over public or transit (recommended).
D. Yes, should use private VIF, but for access to AWS public resources, not the other VPCs.
E. VPC peering wont allow Onprem to access other VPCs via peering.

Best Answer is DX-Gateway AND Public VIF (A and D). However they're both wrong.

https://docs.aws.amazon.com/directconnect/latest/UserGuide/direct-connect-gateways-intro.html
upvoted 25 times

GabrielShiao 5 months, 2 weeks ago

Vote D.
You can access the AWS public resources if you create a public VIF well. By setting the AWS site-to-set VPN, one of AWS's public resources, you can leverage this VPN to connect to the multiple VPC accordingly.
upvoted 1 times

Roontha Highly Voted 2 years, 1 month ago

Answer : A, D
https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-site-to-site-vpn.html
upvoted 12 times

Zac15 Most Recent 5 months ago

Selected Answer: AD

https://docs.aws.amazon.com/whitepapers/latest/aws-direct-connect-for-amazon-connect/virtual-interfaces-vif.html
upvoted 1 times

gfhbox0083 11 months, 3 weeks ago

Selected Answer: AD

A, D for sure.
Must have access to AWS public services.

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: AD

A and D

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: AD

its AD

upvoted 1 times

SkyZeroZx 1 year, 12 months ago

Selected Answer: AD

Answer : A, D

https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-site-to-site-vpn.html

upvoted 1 times

pupsik 2 years ago

Selected Answer: AD

got to use Public VIN in order to connect to AWS Services via Direct Connect.

upvoted 2 times

easytoo 2 years ago

a-d-a-d-a-d-a-d

upvoted 1 times

Jesuisleon 2 years ago

Agree Roontha.

For E, "Create a private VIF with the existing Direct Connect connection to connect to the peered VPCs" is wrong. private VIF can only connect to the vpc which is in the same region with direct connection, you can't extend private VIF to the VPCs in other 2 regions.

upvoted 5 times

rbm2023 2 years, 1 month ago

Selected Answer: AD

agree with A and D tks to Roontha

upvoted 3 times

andreitugui 2 years, 1 month ago

Selected Answer: AD

Answer is A,D

upvoted 1 times

Question #199 *Topic 1*

A company is using an organization in AWS Organizations to manage hundreds of AWS accounts. A solutions architect is working on a solution to provide baseline protection for the Open Web Application Security Project (OWASP) top 10 web application vulnerabilities. The solutions architect is using AWS WAF for all existing and new Amazon CloudFront distributions that are deployed within the organization.

Which combination of steps should the solutions architect take to provide the baseline protection? (Choose three.)

A. Enable AWS Config in all accounts

B. Enable Amazon GuardDuty in all accounts

C. Enable all features for the organization

D. Use AWS Firewall Manager to deploy AWS WAF rules in all accounts for all CloudFront distributions

E. Use AWS Shield Advanced to deploy AWS WAF rules in all accounts for all CloudFront distributions

F. Use AWS Security Hub to deploy AWS WAF rules in all accounts for all CloudFront distributions

Suggested Answer: *CDF*

Community vote distribution

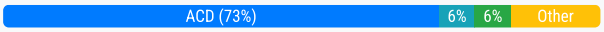

Roontha Highly Voted 2 years, 1 month ago

My Answer A,C,D

https://aws.amazon.com/blogs/security/using-aws-firewall-manager-and-waf-to-protect-your-web-applications-with-master-rules-and-application-specific-rules/

can someone post the link if you feel my answer is incorrect

upvoted 18 times

ShinLi 2 years, 1 month ago

why you pickup C? why we need enable all the features?

upvoted 1 times

Roontha 2 years ago

@ShinLi,

C is must requirement in order leverage AWS Firewall Manager according to aws.

Prerequisites

AWS Firewall Manager has the following prerequisites:

AWS Organizations: Your organization must be using AWS Organizations to manage your accounts, and All Features must be enabled. For more information, see Creating an Organization and Enabling All Features in Your Organization.

A firewall administrator AWS Account: You must designate one of the AWS accounts in your organization as the administrator for AWS Firewall Manager. This gives the account permission to deploy AWS WAF rules across the organization.

AWS Config: You must enable AWS Config for all of the accounts in your organization so that AWS Firewall Manager can detect newly created resources. To enable AWS Config for all of the accounts in your organization, you can use the Enable AWS Config template on the StackSets Sample Templates page. For more information, see Getting Started with AWS Config.

upvoted 22 times

sakibmas Most Recent 9 months, 3 weeks ago

Selected Answer: ACD

AWS Firewall Manager has the following prerequisites:

AWS Organizations: Your organization must be using AWS Organizations to manage your accounts, and All Features must be enabled.

A firewall administrator AWS Account: You must designate one of the AWS accounts in your organization as the administrator for AWS Firewall Manager.

AWS Config: You must enable AWS Config for all of the accounts in your organization so that AWS Firewall Manager can detect newly created

resources.
Reference: https://aws.amazon.com/blogs/security/using-aws-firewall-manager-and-waf-to-protect-your-web-applications-with-master-rules-and-application-specific-rules/
upvoted 2 times

Russs99 1 year, 2 months ago
Selected Answer: ACD
ACD is the correct combination to establish a base line security when deploying within the organization in AWS Organization.
upvoted 1 times

shaaam80 1 year, 6 months ago
Selected Answer: ACD
Answer - ACD
Prerequisites - AWS Config and All Features should be enabled in the organization.
upvoted 2 times

career360guru 1 year, 7 months ago
Selected Answer: ACD
A, C, D
upvoted 1 times

severlight 1 year, 7 months ago
Selected Answer: ACD
AWS config must be enabled in all accounts to identify new resources so AWS Firewall manager works properly
upvoted 3 times

easytoo 1 year, 11 months ago
a-c-d----a-c-d----a-c-d
GuardDuty, Shield Advanced, and Security Hub provide other security capabilities but are not directly related to deploying WAF rules across all accounts and distributions.
upvoted 2 times

NikkyDicky 1 year, 11 months ago
Selected Answer: ACD
its ACD
upvoted 1 times

javitech83 2 years ago
Selected Answer: ACD
D is clear. A and C are needed for D to work

https://aws.amazon.com/es/blogs/security/centrally-manage-aws-waf-api-v2-and-aws-managed-rules-at-scale-with-firewall-manager/#:~:text=Firewall%20Manager%20prerequisites
upvoted 1 times

SkyZeroZx 2 years ago
Selected Answer: ACD
ACD
Link reference : https://aws.amazon.com/es/blogs/security/centrally-manage-aws-waf-api-v2-and-aws-managed-rules-at-scale-with-firewall-manager/#:~:text=Firewall%20Manager%20prerequisites
upvoted 3 times

easytoo 2 years ago
baseline for OWASP = b-d-f
upvoted 1 times

emiliocb4 2 years ago
Selected Answer: ACD
baseline protection vconfiguration.
A to evaluate the configurations of AWS resources
C enabling all features required by Firewall manager
D to enable the waf rules
upvoted 4 times

Jonalb 2 years ago

Selected Answer: ABD

Enable AWS Config in all accounts: AWS Config provides a detailed view of the configuration of AWS resources within an organization. By enabling AWS Config, the solutions architect can track and monitor the configuration of CloudFront distributions and ensure that they adhere to the desired baseline configuration, including AWS WAF settings.

Enable Amazon GuardDuty in all accounts: Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior within AWS accounts. Enabling GuardDuty in all accounts allows for real-time threat detection and alerts related to potential web application vulnerabilities.

upvoted 1 times

SVGoogle89 2 years ago

Prerequisites for using AWS Firewall Manager

Your account must be a member of AWS Organizations

Your AWS account must be a member of an organization in the AWS Organizations service, and the organization must have all features enabled.

Your account must be the AWS Firewall Manager administrator

To configure Firewall Manager policies, your account must be set as the AWS Firewall Manager administrator account, in the Settings pane.

You must have AWS Config enabled for your accounts and Regions

You must enable AWS Config for each of your AWS Organizations member accounts and for each AWS Region that contains resources that you want to protect using AWS Firewall Manager.

upvoted 2 times

Jesuisleon 2 years ago

Selected Answer: ACD

A,C,D is right answer.

Infact My initial choice is B,C,D.

After I rewatch neal Davis' video, GuardDuty is intelligent thread detection service based ML,

it does continuous monitoring for : 1) CloudTrail Management events; 2) CloudTrail S3 Data Events;3)VPC Flow Logs 4) DNS logs. so guardduty is not right in this scenario.

upvoted 3 times

chathur 2 years ago

Selected Answer: ACD

The tutorial is here.

https://aws.amazon.com/blogs/security/centrally-manage-aws-waf-api-v2-and-aws-managed-rules-at-scale-with-firewall-manager/#:~:text=Firewall%20Manager%20prerequisites

upvoted 1 times

Gmail78 1 year, 10 months ago

I assume if you want to secure AWS you need Guard duty enabled, it also interact with AWS WAF: https://aws.amazon.com/blogs/security/how-to-use-amazon-guardduty-and-aws-web-application-firewall-to-automatically-block-suspicious-hosts/

upvoted 1 times

Rajivjain 2 years ago

Selected Answer: BDE

Updating My Vote to BDE

Enabling Amazon GuardDuty will help monitor and detect malicious activity.

Deploying WAF rules via Firewall Manager or Shield Advanced will filter incoming traffic and block common attack patterns. These steps can help protect against many of the most common web application security risks identified by OWASP.

A (Enable AWS Config) is not directly related to providing baseline protection for web applications against OWASP's top 10 vulnerabilities.

C (Enable All Features) is too broad and does not specifically address web application security.

F (Use Security Hub) does not have a native capability to deploy WAF rules at scale.

upvoted 2 times

Question #200 *Topic 1*

A solutions architect has implemented a SAML 2.0 federated identity solution with their company's on-premises identity provider (IdP) to authenticate users' access to the AWS environment. When the solutions architect tests authentication through the federated identity web portal, access to the AWS environment is granted. However, when test users attempt to authenticate through the federated identity web portal, they are not able to access the AWS environment.

Which items should the solutions architect check to ensure identity federation is properly configured? (Choose three.)

A. The IAM user's permissions policy has allowed the use of SAML federation for that user.

B. The IAM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principal.
B. Test users are not in the AWSFederatedUsers group in the company's IdP.

C. The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider, the ARN of the IAM role, and the SAML assertion from IdP.

D. The on-premises IdP's DNS hostname is reachable from the AWS environment VPCs.

E. The company's IdP defines SAML assertions that properly map users or groups. In the company to IAM roles with appropriate permissions.

Suggested Answer: *BDF*

Community vote distribution

Rajivjain Highly Voted 1 year, 7 months ago
Kindly correct the Answers' sequence. A to F
upvoted 24 times

Rajivjain 1 year, 7 months ago
Ref: BDF https://www.examtopics.com/discussions/amazon/view/36355-exam-aws-certified-solutions-architect-professional-topic-1/
upvoted 3 times

andreitugui Highly Voted 1 year, 7 months ago
B) The IAM roles created for the federated users' or federated groups' trust policy have set the SAML provider as the principal.
D) The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider, the ARN of the IAM role, and the SAML assertion from IdP.
F)The company's IdP defines SAML assertions that properly map users or groups. In the company to IAM roles with appropriate permissions.
upvoted 22 times

sarlos Most Recent 7 months, 3 weeks ago
B1,C,E
upvoted 6 times

37b2ab7 1 year, 1 month ago
Selected Answer: BCE
For sure - BCE
upvoted 3 times

severlight 1 year, 1 month ago
Selected Answer: BCE
B1, C, E
upvoted 3 times

dkcloudguru 1 year, 3 months ago
BDF is correct
upvoted 1 times

CloudHandsOn 1 year, 3 months ago
Selected Answer: BCE
B,C, & E was my first choice
upvoted 2 times

Gmail78 1 year, 4 months ago

C- STS AssumerolewithSAML

B1- Define trust policy for IAM assumed by the principal

E - SAML Assertion

upvoted 3 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: BD

BDF is correct

upvoted 1 times

anttan 1 year, 4 months ago

Should be BEF, right?

D. The web portal calls the AWS STS AssumeRoleWithSAML API with the ARN of the SAML provider, the ARN of the IAM role, and the SAML assertion from IdP. This is already being done by the federated identity web portal.

So E) The on-premises IdP's DNS hostname is reachable from the AWS environment VPCs. The on-premises IdP's DNS hostname must be reachable from the AWS environment VPCs. This is because the AWS STS AssumeRoleWithSAML API will need to be able to resolve the DNS hostname of the IdP in order to retrieve the SAML assertion.

upvoted 2 times

breadops 1 year, 5 months ago

Selected Answer: B

BDF is the right answers

upvoted 2 times

ggrodskiy 1 year, 5 months ago

Correct BCE.

upvoted 1 times

Just_Ninja 1 year, 5 months ago

Selected Answer: BD

Admin The Order from the Question is not right.. Answer is BDF!

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: BCE

B (the 1st B, as there are two in this version of question) CE

upvoted 2 times

easytoo 1 year, 6 months ago

it's B-D-F Jeff.

upvoted 2 times

Roontha 1 year, 7 months ago

Answer : B, C, E

upvoted 2 times

Roontha 1 year, 7 months ago

Sorry...it is BDF

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers_saml.html

upvoted 4 times

Question #201 *Topic 1*

A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.

The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.

Which solution will meet these requirements?

A. Deploy an Amazon RDS Proxy layer. In front of the DB instance. Store the connection credentials as a secret in AWS Secrets Manager.

B. Deploy an Amazon RDS Proxy layer in front of the DB instance. Store the connection credentials in AWS Systems Manager Parameter Store

C. Create an Aurora Replica. Store the connection credentials as a secret in AWS Secrets Manager

D. Create an Aurora Replica. Store the connection credentials in AWS Systems Manager Parameter Store.

Suggested Answer: *A*

Community vote distribution

A (100%)

Masonyeoh Highly Voted 2 years, 1 month ago

Selected Answer: A

Using RDS Proxy, you can handle unpredictable surges in database traffic. Otherwise, these surges might cause issues due to oversubscribing connections or creating new connections at a fast rate. RDS Proxy establishes a database connection pool and reuses connections in this pool. This approach avoids the memory and CPU overhead of opening a new database connection each time. To protect the database against oversubscription, you can control the number of database connections that are created.

upvoted 7 times

85b5b55 Most Recent 4 months, 1 week ago

Selected Answer: A

To handle the overloaded connections and keep the secrets in the Amazon Secret Manager.

upvoted 1 times

carpa_jo 1 year, 6 months ago

Selected Answer: A

Use replicas to scale read, this use-case is about writing so C & D are out.

Secret manager offers rotation, parameter store doesn't.

So its A.

upvoted 2 times

duriselvan 1 year, 6 months ago

D. Aurora Replica with Parameter Store:

Pros:

Improves database capacity and reduces load on the primary instance.

Parameter Store provides centralized configuration management.

Cons:

Manually rotating credentials in Parameter Store poses security risks.

upvoted 1 times

helloworldabc 10 months, 1 week ago

just A

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: A

Option A

upvoted 2 times

joleneinthebackyard 1 year, 8 months ago

Selected Answer: A

straight A. love these questions 😁

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: A

easy A

upvoted 1 times

pupsik 2 years ago

Selected Answer: A

Agree with other explanations here.

upvoted 1 times

rbm2023 2 years, 1 month ago

Selected Answer: A

Agree with A

Rotate the keys using Secrets Manager, Param store does not cover it.

RDS Proxy is exactly to solve the issues with overloaded connection because is a connection pool component.

upvoted 3 times

Roontha 2 years, 1 month ago

Answer : A

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

upvoted 4 times

Question #202 *Topic 1*

A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.large Amazon EC2 instances and uses an Amazon RDS for MySQL DB instance. The EC2 instances are in an Auto Scaling group that extends across multiple Availability Zones.

In the event of a disaster, the web application must fail over to the secondary environment with an RPO of 30 seconds and an RTO of 10 minutes.

Which solution will meet these requirements MOST cost-effectively?

A. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Recover the EC2 instances from the latest EC2 backup. Use an Amazon Route 53 geolocation routing policy to automatically fail over to the DR Region in the event of a disaster.

B. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. Run the EC2 instances at the minimum capacity in the DR Region. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. Increase the desired capacity of the Auto Scaling group.

C. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Manually restore the backed-up data on new instances. Use an Amazon Route 53 simple routing policy to automatically fail over to the DR Region in the event of a disaster.

D. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create an Amazon Aurora global database. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. Run the Auto Scaling group of EC2 instances at full capacity in the DR Region. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster.

Suggested Answer: *B*

Community vote distribution

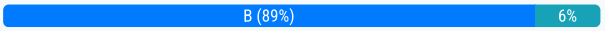

bjexamprep Highly Voted 1 year, 6 months ago

Selected Answer: B

Bad question design. EC2 is in ASG, which means the application part is stateless, so no need to backup or replicate. Only database need replication.

upvoted 7 times

Jonalb Highly Voted 2 years ago

Selected Answer: B

Explanation:

Option B leverages infrastructure as code (IaC) to provision the necessary infrastructure in the DR Region, which allows for automated and repeatable deployments.

Creating a cross-Region read replica for the Amazon RDS DB instance ensures that the database is replicated and available in the DR Region.

AWS Elastic Disaster Recovery can be used to continuously replicate the EC2 instances from the primary Region to the DR Region, ensuring up-to-date copies of the application.

Running the EC2 instances at the minimum capacity in the DR Region helps reduce costs, as resources are only utilized when failover occurs.

Using an Amazon Route 53 failover routing policy allows for automatic failover to the DR Region in the event of a disaster, minimizing downtime.

Increasing the desired capacity of the Auto Scaling group ensures that sufficient resources are available in the DR Region to handle the workload during failover.

upvoted 5 times

FZA24 Most Recent 7 months, 2 weeks ago

Selected Answer: B

RPO seconds, RTO minutes => warm standby

warm standby => always running but smaller

always running but smaller => B. Run the EC2 instances at the minimum capacity in the DR Region

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B most cost effective for RTO=10 min and RPO=30 min.

upvoted 3 times

career360guru 1 year, 7 months ago

RPO=30 sec

upvoted 2 times

Pupu86 1 year, 7 months ago

Selected Answer: B

RPO of 30 seconds can be achieved with Elastic disaster recovery for continuous EC2 instance replication, while DB read replica can be promoted to primary within 30 seconds

upvoted 3 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: B

Close between B & D but Max out ASG is tie-breaker

upvoted 3 times

softarts 1 year, 10 months ago

Selected Answer: D

I think (D) only aurora global database can meet RPO 30 seconds? although B is cost-effective

upvoted 2 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

B for sure

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

A) Not seems for my , posible backup

B) Active Pasive

C) Backup

D) Active Active

Then B is correct in this case

upvoted 3 times

Jackhemo 2 years ago

olabiba.ai said B.

upvoted 1 times

Moallal 2 years ago

Selected Answer: A

Do the math, option A is 5.55 days.

upvoted 1 times

Snape 2 years, 1 month ago

Selected Answer: B

A Wrong - I have stopped reading after 'create cron' , Same goes with C.

D Wrong - Running ASG at full capacity in the DR is not cost efficient

upvoted 5 times

rbm2023 2 years, 1 month ago

i think i agree with option B, initially chosen D

the problem is that we need a cost effective solution and based on the following the global database might be more expensive and the fact the RDS cross region replication may cover the RTO of 10 minutes.

quick compare on global database and cross region replication

RDS Cross Region Replication - You will accrue charges for data transfer between Amazon EC2 and Amazon RDS across Regions, charged on both sides of the transfer ($0.02/GB out)

Aurora Global Database - you pay for replicated write I/O operations between the primary Region and each secondary Region. The number of replicated write I/O operations to each secondary Region is the same as the number of in-Region write I/O operations performed by the primary Region Replicated Write I/Os $0.20 per million replicated write I/Os

upvoted 2 times

andreitugui 2 years, 1 month ago

Selected Answer: B

I would go with B as 10minutes RTO allows for scale up the ASG size. Also read replica is cheaper and can be promoted to primary. Also aurora replication to read replica is usually much less than 100 milliseconds after the primary writes operation which will be enough fot the RPO of 30 seconds.

upvoted 1 times

dbaroger 2 years, 1 month ago

Selected Answer: B

Cost efective = B

upvoted 2 times

AMEJack 2 years, 1 month ago

Selected Answer: B

Agree with B

upvoted 1 times

Masonyeoh 2 years, 1 month ago

Selected Answer: C

save the running EC2 cost. Only bring up when needed

upvoted 1 times

Roontha 2 years, 1 month ago

but the question is saying "web application must fail over to the secondary environment with an RPO of 30 seconds and an RTO of 10 minutes"
How RPO/RTO can be achieved with bare minimum EC2 is up and running in DR site.

Can you paste the link/reading to justify your answer.
Thanks

upvoted 3 times

Question #203 *Topic 1*

A company is planning a one-time migration of an on-premises MySQL database to Amazon Aurora MySQL in the us-east-1 Region. The company's current internet connection has limited bandwidth. The on-premises MySQL database is 60 TB in size. The company estimates that it will take a month to transfer the data to AWS over the current internet connection. The company needs a migration solution that will migrate the database more quickly.

Which solution will migrate the database in the LEAST amount of time?

A. Request a 1 Gbps AWS Direct Connect connection between the on-premises data center and AWS. Use AWS Database Migration Service (AWS DMS) to migrate the on-premises MySQL database to Aurora MySQL.

B. Use AWS DataSync with the current internet connection to accelerate the data transfer between the on-premises data center and AWS. Use AWS Application Migration Service to migrate the on-premises MySQL database to Aurora MySQL.

C. Order an AWS Snowball Edge device. Load the data into an Amazon S3 bucket by using the S3 interface. Use AWS Database Migration Service (AWS DMS) to migrate the data from Amazon S3 to Aurora MySQL.

D. Order an AWS Snowball device. Load the data into an Amazon S3 bucket by using the S3 Adapter for Snowball. Use AWS Application Migration Service to migrate the data from Amazon S3 to Aurora MySQL.

Suggested Answer: *C*

Community vote distribution

C (97%)

F_Eldin Highly Voted 1 year, 7 months ago

Selected Answer: C

Why Not D:

1- C=SnowBall Edge, D=SnowBall Device.

The basic difference between Snowball and Snowball Edge is the capacity they provide. Snowball provides a total of 50 TB or 80 TB, out of which 42 TB or 72 TB is available, while Amazon Snowball Edge provides 100 TB, out of which 83 TB is available.

2- C=AWS Database Migration . D=Application Migration Service,

Application Migration Service simplifies, expedites, and reduces the cost of migrating and modernizing applications. Not for Database

upvoted 27 times

TonytheTiger Most Recent 8 months, 3 weeks ago

Selected Answer: C

Option C : How To

https://aws.amazon.com/blogs/storage/enable-large-scale-database-migrations-with-aws-dms-and-aws-snowball/

upvoted 2 times

Maygam 11 months, 3 weeks ago

Selected Answer: C

AWS Snowball and Snowball Edge refers the same thing. From the Snowball FAQ "AWS Snowball is a service that provides secure, rugged devices, so you can bring AWS computing and storage capabilities to your edge environments, and transfer data into and out of AWS. Those rugged devices are commonly referred to as AWS Snowball or AWS Snowball Edge devices. ". Between C and D, it's C using Snowball edge with AWS DMS.

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: C

Option C - Direct connection would take 1 month

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: C

Basic Snowball edge / DMS use case

upvoted 1 times

Moallal 1 year, 6 months ago

Do the math, option A is 5.55 days. It's A
upvoted 2 times

breadops 1 year, 5 months ago
It can take months to provision a DX connection, its not A.
upvoted 2 times

Jackhemo 1 year, 6 months ago
it takes ages to order a 1G circuit.
upvoted 2 times

covabix879 1 year, 2 months ago
Keyword is one-time migration. In addition to time it takes to deliver, it will be huge waste for one-time task.
upvoted 2 times

andreitugui 1 year, 7 months ago
Selected Answer: C
First of all a snowball solution is required for one time migration will focus in C & D.
Now since we are looking o migrate a database, DMS is needed also Snowball edge can accommodate the 60TB of data as the capacity limit it 80TB.
D is wrong by mentioning Application Migration service to migrate a database.

So correct answer is C). Order an AWS Snowball Edge device. Load the data into an Amazon S3 bucket by using the S3 interface. Use AWS Database Migration Service (AWS DMS) to migrate the data from Amazon S3 to Aurora MySQL.
upvoted 4 times

rbm2023 1 year, 7 months ago
Selected Answer: C
I agree with option C.
Option D does not seem ideal because mentions Application Migration Service, also the snowball is more required for petabyte scale data migration while edge seems to be a better fit.
upvoted 1 times

dbaroger 1 year, 7 months ago
Selected Answer: D
D better cost than C and it does the same for S3. Need adpter too
upvoted 1 times

Roontha 1 year, 7 months ago
Answer : C (Key words : Limited bandwidth + DB migration should be done quickly)

if there no DB migration, we can go with B
upvoted 2 times

Question #204 *Topic 1*

A company has an application in the AWS Cloud. The application runs on a fleet of 20 Amazon EC2 instances. The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes.

The company must maintain backups in a separate AWS Region. The company must be able to recover the EC2 instances and their configuration within 1 business day, with loss of no more than 1 day's worth of data. The company has limited staff and needs a backup solution that optimizes operational efficiency and cost. The company already has created an AWS CloudFormation template that can deploy the required network configuration in a secondary Region.

Which solution will meet these requirements?

A. Create a second CloudFormation template that can recreate the EC2 instances in the secondary Region. Run daily multivolume snapshots by using AWS Systems Manager Automation runbooks. Copy the snapshots to the secondary Region. In the event of a failure launch the CloudFormation templates, restore the EBS volumes from snapshots, and transfer usage to the secondary Region.

B. Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes. In the event of a failure, launch the CloudFormation template and use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region.

C. Use AWS Backup to create a scheduled daily backup plan for the EC2 instances. Configure the backup task to copy the backups to a vault in the secondary Region. In the event of a failure, launch the CloudFormation template, restore the instance volumes and configurations from the backup vault, and transfer usage to the secondary Region.

D. Deploy EC2 instances of the same size and configuration to the secondary Region. Configure AWS DataSync daily to copy data from the primary Region to the secondary Region. In the event of a failure, launch the CloudFormation template and transfer usage to the secondary Region.

Suggested Answer: *C*

Community vote distribution

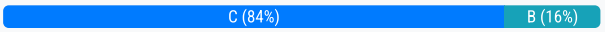

andreitugui Highly Voted 2 years, 1 month ago

Selected Answer: C

Correct is C. For those voting with B, you missed the Instance configuration part. DLM will only backup the EBS volume not the instance settings also. AWS backup will backup ebs & instance settings.

Option C, using AWS Backup, provides a centralized and cost-effective solution for managing backups across multiple services, including EC2 instances. By creating a scheduled daily backup plan for the EC2 instances, AWS Backup ensures regular backups are taken. The backups can be configured to be stored in a vault in the secondary Region, fulfilling the requirement of maintaining backups in a separate Region.
The EC2 instance volumes and configurations can then be restored from the backup vault using AWS Backup's restore capabilities. This allows for the recovery of EC2 instances and their configurations within the required timeframe of 1 business day, with a maximum data loss of 1 day's worth.

upvoted 20 times

Roontha 2 years ago

Answer is B.

https://aws.amazon.com/ebs/data-lifecycle-manager/

It has aws sponsored video which stated clearly can take EBS backed AMIs with AWS DLM

upvoted 1 times

helloworldabc 10 months, 1 week ago

just C

upvoted 1 times

Just_Ninja 1 year, 11 months ago

B is Wrong!
Why? They must!! So that means Compliance is important. AWS Backup is a service for Compliance and Goverment Targets. C Match

upvoted 1 times

sergza888 Most Recent 4 months, 1 week ago

Selected Answer: A

With these RTO/RPO WE don't need to backup entire EC2 especially for cost efficiency. We Only need to maintain CF In another region as well as EBS Backups. System Manager allows you to script and execute backup and copy it to another region instead of DL

upvoted 1 times

chris_spencer 8 months, 2 weeks ago

Selected Answer: C

C because of this one. "The company has limited staff and needs a backup solution that optimizes operational efficiency and cost." AWS Backup really optimizes your backup solution. We backup everthing now with AWS Backup. B works too but it more complicated. The restore from AWS Backup is nearly a no brainer

upvoted 1 times

gfhbox0083 11 months, 3 weeks ago

Selected Answer: C

C, for sure.

Use AWS Backup.

DLM itself does not directly support restore operations.

upvoted 1 times

saggy4 1 year, 4 months ago

Correct Answer is C.

Why not B, DLM can only take backup on restore. The options says using DLM restore the volumes.

upvoted 1 times

saggy4 1 year, 4 months ago

I meant DLM cannot restore so the option B is wrong.

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C

upvoted 2 times

severlight 1 year, 7 months ago

Selected Answer: C

Because AWS Back ups supports restore and DLM doesn't

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: B

B

The explanation here fits the use-case

https://aws.amazon.com/blogs/storage/automating-amazon-ebs-snapshot-and-ami-management-using-amazon-dlm/

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: C

C

B would be ok, if DLM supported restore. it doesn't

upvoted 2 times

javitech83 2 years ago

Selected Answer: C

I think correct is C. AWS Backup is easier and perfectly fits the scenario

upvoted 1 times

Maria2023 2 years ago

Selected Answer: C

B says "Use Amazon DLM to restore the EBS volumes and transfer usage to the secondary Region" - just tested it and could not find any option for DLM to restore volumes, think the snapshots are managed the usual way.

upvoted 2 times

easytoo 2 years ago

C-C-C-C-C-C-C-C-C-C

upvoted 1 times

Jonalb 2 years ago

Selected Answer: B

Its B!!!!!!!!!!!!!!!!

upvoted 1 times

clownfishman 2 years ago

Why not A?

upvoted 3 times

Jesuisleon 2 years ago

Selected Answer: B

I prefer B to C as this sentence "The EC2 instances are persistent and store data on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes", in this question, there is no database mentioned, I assume all persistent data is in EBS, so no need to backup ec2 instances, you can directly startup ec2 instance by cloudformation and load backuped ebs.

upvoted 2 times

rbm2023 2 years, 1 month ago

Selected Answer: C

AWS Backup is more cost effective so I would chose C as well. The DLM option B, does not contemplate the back up in another region as far as I could see.

upvoted 2 times

Jesuisleon 2 years ago

DLM can copy snapshots to another region, see https://aws.amazon.com/about-aws/whats-new/2019/12/amazon-data-lifecycle-manager-enables-automation-snapshot-copy-via-policies/

upvoted 1 times

F_Eldin 2 years, 1 month ago

Selected Answer: B

AWS Backup is a latter service which tries to simplify the challenge of administering a backup in each service individually.

However AWS Lifecycle Manager originally only made EBS snapshots but has been expanded to create AMIs. I don't believe AWS Backup can trigger AMI creation.

upvoted 1 times

andreitugui 2 years, 1 month ago

But B mentions only EBS snapshots (Use Amazon Data Lifecycle Manager (Amazon DLM) to create daily multivolume snapshots of the EBS volumes)! Does not say anything about AMI's.

So IMO the answer is C

upvoted 1 times

Question #205 *Topic 1*

A company is designing a new website that hosts static content. The website will give users the ability to upload and download large files. According to company requirements, all data must be encrypted in transit and at rest. A solutions architect is building the solution by using Amazon S3 and Amazon CloudFront.

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

A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.

B. Add a policy attribute of "aws:SecureTransport": "true" for read and write operations in the S3 ACLs.

C. Create a bucket policy that denies any unencrypted operations in the S3 bucket that the web application uses.

D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE-KMS).

E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.

F. Use the RequireSSL option in the creation of presigned URLs for the S3 bucket that the web application uses.

Suggested Answer: *ACE*

Community vote distribution

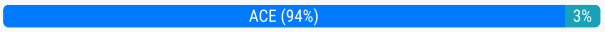

SkyZeroZx Highly Voted 1 year, 12 months ago

Selected Answer: ACE

Answer : ACE

A) SSE S3 sounds good encript in rest data

B) sounds good until say in ACLs is incorrect

C) Bucket Policy avoid upload unencrypted is correct sounds good

D) CloudFront with KMS ? why ? not seems

E) HTTP redirect to HTTPS sounds good is clasic this case

F) why ? not seems in this case

upvoted 19 times

Just_Ninja Highly Voted 1 year, 11 months ago

Selected Answer: ACE

ACE.

But A is deprecated :)

because since the 05.01.2023 S3 use automatical atRest encryption for new objekts.

upvoted 6 times

dankositzke 1 year, 4 months ago

Right I would go with CEF for 2024 onwards

upvoted 2 times

khchan123 Most Recent 1 year, 3 months ago

Selected Answer: BCE

BCE

You need B to enforce encryption in transit with S3. Other options cannot do that.

upvoted 1 times

helloworldabc 10 months, 1 week ago

just ACE

upvoted 1 times

Dgix 1 year, 3 months ago

Selected Answer: ACE

This question was obviously formulated before S3 buckets were encrypted by default.

upvoted 1 times

duriselvan 1 year, 6 months ago

A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.
D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE-KMS).
E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.

Here's why these steps are necessary:

A. S3 server-side encryption: This encrypts data in the S3 bucket at rest, ensuring data confidentiality even if someone gains unauthorized access to the bucket.
D. CloudFront SSE-KMS: This encrypts data in transit between CloudFront and the client, ensuring data confidentiality when users upload and download files.
E. HTTP to HTTPS redirect: This ensures all communication between the client and CloudFront occurs over HTTPS, encrypting data in transit and preventing eavesdropping.
upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: ACE

Options A, C , E
upvoted 1 times

task_7 1 year, 9 months ago

Selected Answer: ADE

A. Turn on S3 server-side encryption for the S3 bucket that the web application uses.
D. Configure encryption at rest on CloudFront by using server-side encryption with AWS KMS keys (SSE-KMS).
E. Configure redirection of HTTP requests to HTTPS requests in CloudFront.
Data at rest encrypted for Both S3 and Cloudfront
E for data in transit
upvoted 1 times

Simon523 1 year, 10 months ago

Selected Answer: ACE

How to Prevent Uploads of Unencrypted Objects to Amazon S3
https://aws.amazon.com/tw/blogs/security/how-to-prevent-uploads-of-unencrypted-objects-to-amazon-s3/
upvoted 2 times

RotterDam 1 year, 10 months ago

ACE but why not F?
upvoted 1 times

chikorita 1 year, 9 months ago

question nowhere mentions the use of pre-signed URLs
if it was used in this scenario then it could potentially be one of the right answers
upvoted 3 times

kgpoj 10 months, 3 weeks ago

When you have pre-signed urls, you don't even necessarily need cloudFront
upvoted 1 times

Christina666 1 year, 11 months ago

Selected Answer: ACE

we don't have a "encrytion at rest" for cloudfront in the console
upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: ACE

A and C are a bit redundant. I'd pick D instead of C, but for ACL reference
upvoted 1 times

easytoo 2 years ago

a-d-e a-d-e a-d-e
upvoted 2 times

chathur 2 years ago

Selected Answer: ACE

Source: https://repost.aws/knowledge-center/s3-bucket-policy-for-config-rule

B is wrong as "aws:SecureTransport": "true" does not deny 'http' traffic

upvoted 1 times

consultornetwork 2 years ago

Why not B?

upvoted 2 times

Jesuisleon 2 years ago

you should add "aws:SecureTransport": "true" in the S3 bucket policy not S3 ACL.
see https://stackoverflow.com/questions/47815526/s3-bucket-policy-vs-access-control-list

and " We recommend allowing only encrypted connections over HTTPS (TLS) by using the aws:SecureTransport condition in your Amazon S3 bucket policies" from https://docs.aws.amazon.com/AmazonS3/latest/userguide/security-best-practices.html

upvoted 6 times

BabaP 2 years ago

Because C does just that

upvoted 1 times

chathur 2 years ago

https://repost.aws/knowledge-center/s3-bucket-policy-for-config-rule
it is not enough

upvoted 1 times

andreitugui 2 years, 1 month ago

Selected Answer: ACE

I will go with ACE

upvoted 2 times

Roontha 2 years, 1 month ago

Answer : ACE

upvoted 4 times

Question #206 *Topic 1*

A company is implementing a serverless architecture by using AWS Lambda functions that need to access a Microsoft SQL Server DB instance on Amazon RDS. The company has separate environments for development and production, including a clone of the database system.

The company's developers are allowed to access the credentials for the development database. However, the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access. This key must be rotated on a regular basis.

What should a solutions architect do in the production environment to meet these requirements?

A. Store the database credentials in AWS Systems Manager Parameter Store by using a SecureString parameter that is encrypted by an AWS Key Management Service (AWS KMS) customer managed key. Attach a role to each Lambda function to provide access to the SecureString parameter. Restrict access to the SecureString parameter and the customer managed key so that only the IT security team can access the parameter and the key.

B. Encrypt the database credentials by using the AWS Key Management Service (AWS KMS) default Lambda key. Store the credentials in the environment variables of each Lambda function. Load the credentials from the environment variables in the Lambda code. Restrict access to the KMS key so that only the IT security team can access the key.

C. Store the database credentials in the environment variables of each Lambda function. Encrypt the environment variables by using an AWS Key Management Service (AWS KMS) customer managed key. Restrict access to the customer managed key so that only the IT security team can access the key.

D. Store the database credentials in AWS Secrets Manager as a secret that is associated with an AWS Key Management Service (AWS KMS) customer managed key. Attach a role to each Lambda function to provide access to the secret. Restrict access to the secret and the customer managed key so that only the IT security team can access the secret and the key.

Suggested Answer: *D*

Community vote distribution

Snape Highly Voted 2 years, 1 month ago

Selected Answer: D

Answer : D

Rotation = Secret Manager (and Not Parameter store)

upvoted 13 times

Jassybanga Most Recent 10 months ago

Answer should be A , As we are talking of encryption Key rotation by customer IT key responisble person and not the database credential rotation

upvoted 2 times

AA001 10 months, 1 week ago

Selected Answer: D

To use parameters from Parameter Store in AWS Lambda functions without using an SDK, you can use the AWS Parameters and Secrets Lambda Extension.

To use parameters in a Lambda function without the Lambda extension, you must configure your Lambda function to receive configuration updates by integrating with the GetParameter API action for Parameter Store.

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: D

Option D

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: D

its a D

upvoted 1 times

javitech83 2 years ago

Selected Answer: D

Keys is DB credentials rotation

upvoted 2 times

easytoo 2 years ago

d-d-d-d-dd-d-dd-d-d-d

upvoted 1 times

Jackhemo 2 years ago

Selected Answer: A

From olabiba.ai

"Based on the requirements of resolving scaling issues and minimizing licensing costs, the most cost-effective solution would be option A: Deploy Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer for the web tier and for the application tier. Use Amazon Aurora PostgreSQL with Babelfish turned on to replatform the SQL Server database."

upvoted 1 times

Just_Ninja 1 year, 11 months ago

Nice description, but A is Wrong. Parameter Store is not the best practice for Secrets based on AWS Well Architecting Framework

upvoted 2 times

Jackhemo 2 years ago

Answer is D. This is for the next question.

upvoted 2 times

rbm2023 2 years, 1 month ago

Selected Answer: A

I think the answer is A the requirement is to rotate the KEY and not the password, looks like this question was created to make us chose option D. Option A stores the password in the Param Store encrypting it with KMS which is the requirement “the credentials for the production database must be encrypted with a key that only members of the IT security team's IAM user group can access.”

https://docs.aws.amazon.com/systems-manager/latest/userguide/ps-integration-lambda-extensions.html

Check the Authentication section.

upvoted 4 times

F_Eldin 2 years ago

A does not satisfy the requirement "This key must be rotated on a regular basis."

upvoted 3 times

kejam 1 year, 5 months ago

Agreed. Requirement is to rotate the Key. KMS CMKs can be rotated:

https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html

upvoted 1 times

andreitugui 2 years, 1 month ago

Selected Answer: D

Answering D

upvoted 1 times

Masonyeoh 2 years, 1 month ago

Selected Answer: D

D, Secret Manager is the accurate solution

upvoted 1 times

Roontha 2 years, 1 month ago

Answer : D

Keys is DB credentials rotation

upvoted 1 times

Question #207 *Topic 1*

An online retail company is migrating its legacy on-premises .NET application to AWS. The application runs on load-balanced frontend web servers, load-balanced application servers, and a Microsoft SQL Server database.

The company wants to use AWS managed services where possible and does not want to rewrite the application. A solutions architect needs to implement a solution to resolve scaling issues and minimize licensing costs as the application scales.

Which solution will meet these requirements MOST cost-effectively?

A. Deploy Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer for the web tier and for the application tier. Use Amazon Aurora PostgreSQL with Babelfish turned on to replatform the SQL Server database.

B. Create images of all the servers by using AWS Database Migration Service (AWS DMS). Deploy Amazon EC2 instances that are based on the on-premises imports. Deploy the instances in an Auto Scaling group behind a Network Load Balancer for the web tier and for the application tier. Use Amazon DynamoDB as the database tier.

C. Containerize the web frontend tier and the application tier. Provision an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Create an Auto Scaling group behind a Network Load Balancer for the web tier and for the application tier. Use Amazon RDS for SQL Server to host the database.

D. Separate the application functions into AWS Lambda functions. Use Amazon API Gateway for the web frontend tier and the application tier. Migrate the data to Amazon S3. Use Amazon Athena to query the data.

Suggested Answer: *A*

Community vote distribution

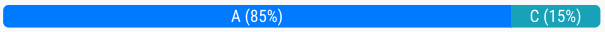

bjexamprep Highly Voted 1 year, 6 months ago

Selected Answer: A

"does not want to rewrite the application. " leaves the possible answer between A and C, cause B and D will force the application team to rewrite the data access part of the application.
C is using EKS, which makes AutoScalingGroup is not required. ASG scales instances. ASG doesn't scale PODs in EKS.
Babelfish is the key point in this question. "Babelfish for Aurora PostgreSQL is a new capability for Amazon Aurora PostgreSQL-Compatible Edition that enables Aurora to understand commands from applications written for Microsoft SQL Server."

upvoted 13 times

F_Eldin Highly Voted 2 years ago

Selected Answer: A

There is no good solution here. A is just forcing that company to use AWS services as "MOST cost-effectively" alternative. Practically Bablefish has bad reviews, companies prefer to migrate SQL-Server as-is.

upvoted 6 times

85b5b55 Most Recent 1 month, 3 weeks ago

Selected Answer: C

C - Fully managed Service, cost-effective and doesn't want to rewrite, those points pushed me to select C only.

upvoted 1 times

SIJUTHOMASP 6 months, 1 week ago

Selected Answer: A

The key is 'minimise licensing cost' so, option A is the best because it can radically cut down the SQL Server licensing cost by putting it to Aurora PostgreSQL. Option C has equivalent licensing cost since it is SQL RDS.

upvoted 2 times

0b43291 7 months, 3 weeks ago

Selected Answer: C

Option C is the most cost-effective solution as it leverages containerization with Amazon EKS, Auto Scaling groups with Network Load Balancers, and Amazon RDS for SQL Server. This approach allows for efficient scaling, resource utilization, and minimizes licensing costs without requiring significant application changes.

Containerizing the web and application tiers enables portability and scalability. Amazon EKS provides a fully managed Kubernetes service, reducing operational overhead. Auto Scaling groups and Network Load Balancers enable automatic scaling based on demand. Amazon RDS for SQL Server offers a fully managed database service with various licensing models, including BYOL, to optimize costs as the application scales.

The other options have drawbacks, such as requiring replatforming the database (Option A), significant application changes (Option B), or a complete rewrite (Option D), which goes against the requirements.

upvoted 1 times

8693a49 11 months ago

Selected Answer: C

All answers are wrong. A is not using managed services where possible (EKS would be better than EC2 and can run windows) and on C you can't have Auto Scaling group for EKS. Realistically C is the better option if scaled with Karpenter, etc.

upvoted 1 times

helloworldabc 10 months, 1 week ago

just A

upvoted 1 times

BrijMohan08 1 year, 1 month ago

Selected Answer: C

Key here is AWS Managed = EKS

A. Says both Web tier and Application tier is behind ALB, which is not secure.

A good design should have web tier behind ALB, and application tier behind NLB

upvoted 1 times

TonytheTiger 1 year, 3 months ago

Selected Answer: A

Option A: Babelfish for Aurora PostgreSQL is a capability for Amazon Aurora PostgreSQL-Compatible Edition developed using the PostgreSQL extension framework that enables Aurora to understand commands from applications written for Microsoft SQL Server. Babelfish for Aurora PostgreSQL understands T-SQL, Microsoft SQL Server's SQL dialect, and supports

https://aws.amazon.com/blogs/database/run-sql-server-reporting-services-reports-against-babelfish-for-aurora-postgresql/

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: A

Option A

upvoted 1 times

Pupu86 1 year, 7 months ago

Selected Answer: C

As much as I would like to choose A but the question request for lift and shift approach rather than a replatform

upvoted 2 times

enk 1 year, 7 months ago

Selected Answer: C

I vote C. Babelfish - another layer to keep an eye on. Is it really going to translate all SQL app calls perfectly, or will they need tuning?

upvoted 1 times

kjcncjek 1 year, 10 months ago

why not C

upvoted 1 times

Mikado211 1 year, 7 months ago

C would be probably the most realistic way a team work to engage such case regarding to the choices we have. However Babelfish is a tool made to execute Microsoft SQL on a postgreSQL server. In practice Babelfish is a toy and should not be used for a real strong usage since the database engine is the last thing you want to play with. Still, people who answered A have followed the theory, and it's probably the expected answer here.

upvoted 3 times

chikorita 1 year, 10 months ago

A : the best of the worst

upvoted 3 times

ggrodskiy 1 year, 11 months ago

Correct A.

upvoted 1 times

YodaMaster 1 year, 11 months ago

Selected Answer: A

A. The other options sound fishy.

upvoted 5 times

rxhan 1 year, 11 months ago

golden.

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: A

A by elimination

upvoted 2 times

easytoo 2 years ago

a-a-a-a-a-a-a

after much consideration it's the babelfish to the rescue - zaphod beeblebrox ftw

upvoted 1 times

Question #208 *Topic 1*

A software-as-a-service (SaaS) provider exposes APIs through an Application Load Balancer (ALB). The ALB connects to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster that is deployed in the us-east-1 Region. The exposed APIs contain usage of a few non-standard REST methods: LINK, UNLINK, LOCK, and UNLOCK.

Users outside the United States are reporting long and inconsistent response times for these APIs. A solutions architect needs to resolve this problem with a solution that minimizes operational overhead.

Which solution meets these requirements?

A. Add an Amazon CloudFront distribution. Configure the ALB as the origin.

B. Add an Amazon API Gateway edge-optimized API endpoint to expose the APIs. Configure the ALB as the target.

C. Add an accelerator in AWS Global Accelerator. Configure the ALB as the origin.

D. Deploy the APIs to two additional AWS Regions: eu-west-1 and ap-southeast-2. Add latency-based routing records in Amazon Route 53.

Suggested Answer: *C*

Community vote distribution

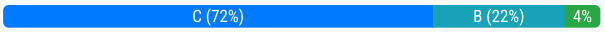

andreitugui Highly Voted 2 years, 1 month ago

Selected Answer: C

AWS Global Accelerator is a service that improves the availability and performance of applications for global users. By adding an accelerator in AWS Global Accelerator and configuring the ALB as the origin, the traffic from users outside the United States will be routed through the Global Accelerator network, which uses the AWS global network infrastructure to optimize the delivery of the application traffic.

upvoted 11 times

sam2ng 10 months ago

as the ALB and EKS are running in one region only which is us-east-1, how does the global accelerator help when traffic comes from other region e.g. EU ?

Also you don't configure origin in global accelerator, you configure endpoint group.

upvoted 1 times

nexus2020 2 years ago

Yes you can, see - https://docs.aws.amazon.com/global-accelerator/latest/dg/introduction-how-it-works.html

--> For standard accelerators, the endpoints are Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses.

upvoted 2 times

ayadmawla Highly Voted 1 year, 6 months ago

Selected Answer: C

imho answer is C. Here is my thinking: There are two issues that we need to consider:

1- Non US Users are reporting long and inconsistent response times for these APIs

2- The APIs are running in EKS and are exposed by the ALB (i.e., not the other way round)

So the issue is about latency not API design.

upvoted 5 times

fabriciollf Most Recent 8 months, 4 weeks ago

Selected Answer: C

C

upvoted 1 times

Helpnosense 1 year ago

Selected Answer: C

Not B. Because Gateway Edge-Optimized API Endpoint improve the performance by caching API responses. But (un)link the call is not supported by API Gateway so the rest will be passed to ALB anyway. So unlikely API gateway will cache and no benefit for performance improvement.

upvoted 2 times

titi_r 1 year, 1 month ago

Selected Answer: C

Answer: C

AWS Global Accelerator is a service in which you create accelerators to improve the performance of your applications for local and global users.
https://docs.aws.amazon.com/global-accelerator/latest/dg/what-is-global-accelerator.html

When you create an ALB or NLB, you can optionally add an accelerator at the same time. Elastic Load Balancing and Global Accelerator work together to transparently add the accelerator for you. The accelerator is created in your account, with the load balancer as an endpoint. Using an accelerator provides static IP addresses and improves the availability and performance of your applications.
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html

upvoted 1 times

titi_r 1 year, 1 month ago

Ans “B” is wrong, because API Gateway does NOT support non-standard REST methods. The supported methods are DELETE, GET, HEAD, OPTIONS, PATCH, POST, PUT, and ANY (which can substitute any of the other 7).
https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-method-settings-method-request.html#setup-method-add-http-method .

upvoted 4 times

dankositzke 1 year, 4 months ago

Selected Answer: D

A: No
B: No, API Gateway doesn’t support LINK, UNLINK, LOCK, UNLOCK.
C: No, GA doesn’t have the concept of “origin” - this is a CloudFront concept.
D: Yes, because this addresses the main concern which is latency.

upvoted 2 times

duriselvan 1 year, 6 months ago

b IS ANS
Minimal operational overhead: API Gateway edge-optimized endpoints offer several advantages:

Reduced latency: They leverage AWS's global network of edge locations, significantly reducing latency for users outside the United States.
Scalability: They automatically scale to handle traffic spikes, eliminating the need for manual intervention.
Security: They offer built-in security features, including access control and throttling, minimizing the need for additional configuration.
Non-standard methods compatibility: API Gateway supports a wide range of HTTP methods, including custom methods like LINK, UNLINK, LOCK, and UNLOCK, ensuring compatibility with the existing APIs.

Ease of configuration: Configuring API Gateway with ALB as the target is straightforward and requires minimal changes to the existing infrastructure.

upvoted 1 times

awsamar 1 year, 6 months ago

Selected Answer: B

Amazon CloudFront primarily supports standard HTTP/HTTPS request methods like GET, POST, PUT, DELETE, HEAD, OPTIONS, and PATCH. It does not natively support non-standard methods such as LINK and UNLINK, LOCK...etc
HOWEVER>>>>>
If you need to use these non-standard methods, you have a couple of options:
Custom Handling with Lambda@Edge
API Gateway Integration: If you require more complex routing and method handling, integrating AWS API Gateway with CloudFront might be a more suitable solution. API Gateway provides robust support for various HTTP methods and can be set up to handle non-standard methods.

Clearly its B

upvoted 3 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C, GA is safest option.

upvoted 1 times

severlight 1 year, 7 months ago

Selected Answer: C

there is no proper use case for API gateway here

upvoted 2 times

joleneinthebackyard 1 year, 8 months ago

Selected Answer: C

A is invalid because cloudFront only support standard Rest Methods

B C D all technically feasible but let's consider "minimized operational overhead" requirement, it's must be C.

upvoted 2 times

chico2023 1 year, 10 months ago

Selected Answer: B

Answer: B

I don't understand why people are choosing GA. I would rather go with option D.

From AWS documentation:

Edge-optimized API endpoint

The default hostname of an API Gateway API that is deployed to the specified Region while using a CloudFront distribution to facilitate client access typically from across AWS Regions. API requests are routed to the nearest CloudFront Point of Presence (POP), which typically improves connection time for geographically diverse clients.

I couldn't find any document mentioning that Edge-optimized API endpoints won't support non-standard REST methods.

upvoted 2 times

chico2023 1 year, 10 months ago

I know we can't trust AI assistants, but take a look at my little chat with:

=== Labiba

Yes, Amazon API Gateway Edge-optimized APIs can handle non-standard REST methods. Edge-optimized APIs are designed to provide low-latency access to your API by using the AWS CloudFront global network. You can set up API methods to handle any HTTP method, including non-standard ones, and configure them to work with your specific requirements and use cases.

=== Bard

Yes, Amazon API Gateway edge-optimized APIs can handle non-standard REST methods. However, there are some limitations.

The non-standard REST method must be supported by the integration that you use for the API method. For example, if you are using a Lambda integration, the Lambda function must be able to handle the non-standard REST method.

upvoted 1 times

chico2023 1 year, 10 months ago

Now, why would I use GA?

I don't know you, but I would use in a situation where I have an application that connects to a database and I need to reduce the latency of my application for users by launching EC2 instances around the world. Note that I can't do that (not that easy, at least) with my RDS DB, so what I do? I use Global Accelerator to speed up communication between my instances in different countries to the database server in a single location, for example.

upvoted 1 times

vn_thanhtung 1 year, 10 months ago

https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-api-endpoint-types.html#api-gateway-api-endpoint-types-edge-optimized:~:text=traffic%20originates%20from.-,Edge%2Doptimized%20API%20endpoints,-An%20edge%2Doptimized

I think can help you, C is answer

upvoted 1 times

Arnaud92 1 year, 11 months ago

Selected Answer: C

Cloudfront cannot handle non standard REST methods. There are Clod front involved behind API Gateway edge-optimized. So only C make sense here

upvoted 4 times

Just_Ninja 1 year, 11 months ago

Selected Answer: B

It only can B...

Here is a AWS entry. https://repost.aws/knowledge-center/api-gateway-cloudfront-distribution

upvoted 3 times

NikkyDicky 1 year, 11 months ago

Selected Answer: C

B would be nice if edge-optimized was supported for HTTP APIs

upvoted 3 times

SandyIndia 2 years ago

Selected Answer: C

By adding an accelerator in AWS Global Accelerator and configuring the ALB as the origin, the traffic to the ALB will be routed through the global network, reducing latency and improving response times for users outside the United States.

This solution minimizes operational overhead as AWS Global Accelerator handles the routing and optimization automatically, without requiring additional infrastructure deployment or configuration changes.

upvoted 3 times

Maria2023 2 years ago

Selected Answer: C

I was also supporting answer B, however just tested API Gateway and it seems that it only supports GET, POST, PUT, PATCH, DELETE, HEAD, and OPTIONS methods. I personally couldn't find a way to create a custom method which is part of the requirement. Please share if you find a way

upvoted 4 times

Question #209 *Topic 1*

A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MQTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named iot.example.com. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.

On several occasions, the amount of data has overloaded the MQTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.

Which solution will meet these requirements?

A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. Use the Auto Scaling group as the target for the ALB. Update the DNS record in Route 53 to an alias record. Point the alias record to the ALB. Use the MQTT broker to store the data.

B. Set up AWS IoT Core to receive the sensor data. Create and configure a custom domain to connect to AWS IoT Core. Update the DNS record in Route 53 to point to the AWS IoT Core Data-ATS endpoint. Configure an AWS IoT rule to store the data.

C. Create a Network Load Balancer (NLB). Set the MQTT broker as the target. Create an AWS Global Accelerator accelerator. Set the NLB as the endpoint for the accelerator. Update the DNS record in Route 53 to a multivalue answer record. Set the Global Accelerator IP addresses as values. Use the MQTT broker to store the data.

D. Set up AWS IoT Greengrass to receive the sensor data. Update the DNS record in Route 53 to point to the AWS IoT Greengrass endpoint. Configure an AWS IoT rule to invoke an AWS Lambda function to store the data.

Suggested Answer: *C*

Community vote distribution

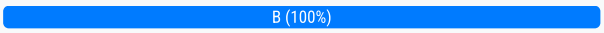

easytoo Highly Voted 2 years ago

b-b-b-b-bb-

Greengrass is typically used for edge computing scenarios and may not be the most suitable solution for addressing MQTT broker reliability and scalability.

upvoted 5 times

junja Most Recent 1 year, 3 months ago

Selected Answer: B

option B

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B

upvoted 1 times

bur4an 1 year, 9 months ago

I think this is repeat question.

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: B

AWS service is the answer.

upvoted 3 times

lferrari 1 year, 10 months ago

Selected Answer: B

IOT core for anything IOT

upvoted 2 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

IOT core for anything IOT

upvoted 4 times

pupsik 2 years ago

Selected Answer: B

Option C doesn't mention required auto-scaling group, hence eliminated.

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

voting for B. IoT Core

upvoted 3 times

Maria2023 2 years ago

Selected Answer: B

Both C and B should work. I suggest AWS wants us to use as many native services as we can, therefore B should be the preferred answer.

upvoted 2 times

Daniel76 7 months, 4 weeks ago

IoT core support availability whereas option c did not mention about auto scaling. With just one instance it might still fail to process when there's a surge in incoming data.

upvoted 1 times

chiaseed 2 years ago

Selected Answer: B

voting for B. IoT Core

upvoted 2 times

nexus2020 2 years ago

Selected Answer: B

IoT core, B

upvoted 1 times

Question #210 *Topic 1*

A company has Linux-based Amazon EC2 instances. Users must access the instances by using SSH with EC2 SSH key pairs. Each machine requires a unique EC2 key pair.

The company wants to implement a key rotation policy that will, upon request, automatically rotate all the EC2 key pairs and keep the keys in a securely encrypted place. The company will accept less than 1 minute of downtime during key rotation.

Which solution will meet these requirements?

A. Store all the keys in AWS Secrets Manager. Define a Secrets Manager rotation schedule to invoke an AWS Lambda function to generate new key pairs. Replace public keys on EC2 instances. Update the private keys in Secrets Manager.

B. Store all the keys in Parameter Store, a capability of AWS Systems Manager, as a string. Define a Systems Manager maintenance window to invoke an AWS Lambda function to generate new key pairs. Replace public keys on EC2 instances. Update the private keys in Parameter Store.

C. Import the EC2 key pairs into AWS Key Management Service (AWS KMS). Configure automatic key rotation for these key pairs. Create an Amazon EventBridge scheduled rule to invoke an AWS Lambda function to initiate the key rotation in AWS KMS.

D. Add all the EC2 instances to Fleet Manager, a capability of AWS Systems Manager. Define a Systems Manager maintenance window to issue a Systems Manager Run Command document to generate new key pairs and to rotate public keys to all the instances in Fleet Manager.

Suggested Answer: *A*

Community vote distribution

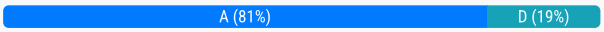

xerxersxu 2 months, 2 weeks ago

Selected Answer: A

https://aws.amazon.com/cn/blogs/security/how-to-use-aws-secrets-manager-securely-store-rotate-ssh-key-pairs/

upvoted 1 times

pk0619 6 months, 1 week ago

Selected Answer: D

SSM RunCommand is the only solution that can actually replace the keys on EC2 instances.

upvoted 1 times

dankositzke 10 months, 2 weeks ago

Selected Answer: A

Not sure why you would need to “invoke an AWS Lambda function to generate new key pairs” when Secrets Manager natively supports automatic key rotation? Anyways, A seems to be the least worst answer.

upvoted 3 times

sat2008 10 months, 1 week ago

Lambda is part of the key creation and rotation see the link

https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-securely-store-rotate-ssh-key-pairs/

upvoted 4 times

Maygam 1 year ago

Selected Answer: A

https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-securely-store-rotate-ssh-key-pairs/

upvoted 3 times

pk0619 6 months, 1 week ago

this is a 5 years old solution, currently the answer should be either B or best D, also Lambda cannot replace the public keys on EC2 instances, you need SSM RunCommand for that.

upvoted 1 times

CProgrammer 1 year ago

@duriselvan ==> How did you arrive at "Automatic key rotation" from "key rotation policy that will, upon request

B. Parameter Store: While Parameter Store can store keys, it's not designed for automated key rotation. It would require manual configuration and orchestration.

C. AWS KMS: KMS is designed for managing encryption keys, not SSH key pairs.
It doesn't support the rotation of SSH key pairs on EC2 instances.
D. Fleet Manager: Fleet Manager, while facilitating management tasks on EC2 instances,
doesn't intrinsically handle key rotation.
It would require integration with other services and custom scripts.
upvoted 1 times

duriselvan 1 year ago
C ans
Automatic key rotation: AWS KMS automatically rotates keys according to the configured schedule, eliminating the need for manual intervention and ensuring timely key updates.

Less than 1 minute downtime: AWS KMS allows for seamless key rotation with minimal downtime. The old key remains active until the new key is generated and propagated, ensuring uninterrupted access to instances.

Secure storage: AWS KMS provides a highly secure and encrypted environment for storing cryptographic keys, exceeding the security offered by Parameter Store.

Lambda function integration: The EventBridge rule can trigger a Lambda function to perform additional tasks during key rotation, such as updating user access controls or notifying administrators.
upvoted 3 times

Jay_2pt0_1 1 year, 1 month ago
Torn between A and D. I don't like the do-it-yourself nature (Lambda) of A, but I understand what everyone is saying about the unique key requirement, which would seem to imply that D is wrong. Don't know tbh.
upvoted 1 times

career360guru 1 year, 1 month ago
Selected Answer: A
Option A
upvoted 1 times

severlight 1 year, 1 month ago
Selected Answer: A
A will work, don't overthink, you can request secret rotation in the Secrets manager, and secrets will be stored in a safe place
upvoted 2 times

Sab 1 year, 2 months ago
Selected Answer: A
D is best option if we need to rotate for all Ec2 with same key pair. Since each EC2 to have a different Key pair, will be better to store in Secrets Manager and have that rotated using lambda.
upvoted 1 times

wahaha2023 1 year, 4 months ago
Selected Answer: A
I think the Systems Manager maintenance window is to perform some potentially disruptive actions, which means the duration of the window is equal to system downtime. and I check the white paper, I seems the duration of system maintenance window should be longer than 1 hour.
upvoted 3 times

chico2023 1 year, 4 months ago
Selected Answer: D
Seriously, all. While it can be done in A, it's better to do that with D. Here is why:

Question says:
"A company has Linux-based Amazon EC2 instances." and "Each machine requires a unique EC2 key pair."

We might be talking about thousands of EC2 instances. But let's continue. Option A says:
"Store all the keys in AWS Secrets Manager." which is OK, you can store up to 500,000 apparently but, seriously, think about. Instances are generated and deleted all the time. This would be cumbersome, even if you do that programmatically. Not convinced? Let me continue.
upvoted 1 times

vn_thanhtung 1 year, 4 months ago
With D how to "keep the keys in a securely encrypted place" ? Should be A

upvoted 1 times

chico2023 1 year, 4 months ago

Same option A, says the following: "Define a Secrets Manager rotation schedule to invoke an AWS Lambda function to generate new key pairs. Replace public keys on EC2 instances."

Now, this is A lot, but how are we going to replace the public keys on EC2 instances? Answer doesn't say.

Finally, for those who are supporting their answer on an AWS blog showing how to use SM to rotate SSH key to manage servers, pay attention to this part: "A secret is created in AWS Secrets Manager. The secret holds the SSH keypair that the master node will use to connect to the other nodes in the cluster."

Their design is "one to many", that is not part of what question says, and I would like to remind you "Each machine requires a unique EC2 key pair."

upvoted 1 times

wahaha2023 1 year, 4 months ago

I am curious about how we can define a 1-minute Systems Manager maintenance window.

upvoted 2 times

easytoo 1 year, 5 months ago

a-a-a-a-a-a-a-a

upvoted 1 times

Just_Ninja 1 year, 5 months ago

Selected Answer: A

A: Based on the Well Architecting Framework for best Practices and that tutorial :) https://aws.amazon.com/de/blogs/security/how-to-use-aws-secrets-manager-securely-store-rotate-ssh-key-pairs/

upvoted 1 times

nicecurls 1 year, 5 months ago

Selected Answer: D

Why A? Select D

upvoted 2 times

Just_Ninja 1 year, 5 months ago

D is wrong, Parameter Store is a good practice to store Parameters but not the Secrets. I know you can use KMS to encrypt the Parameters, but you need a secure store für Secrets and here we have for exmaple the secret manager with FIPS 140-2 Standard.

upvoted 2 times

YodaMaster 1 year, 5 months ago

Selected Answer: A

going with A

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

as someone pointed out D breaks the requirement for unique keys

upvoted 1 times

Question #211 *Topic 1*

A company wants to migrate to AWS. The company is running thousands of VMs in a VMware ESXi environment. The company has no configuration management database and has little knowledge about the utilization of the VMware portfolio.

A solutions architect must provide the company with an accurate inventory so that the company can plan for a cost-effective migration.

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

A. Use AWS Systems Manager Patch Manager to deploy Migration Evaluator to each VM. Review the collected data in Amazon QuickSight. Identify servers that have high utilization. Remove the servers that have high utilization from the migration list. Import the data to AWS Migration Hub.

B. Export the VMware portfolio to a .csv file. Check the disk utilization for each server. Remove servers that have high utilization. Export the data to AWS Application Migration Service. Use AWS Server Migration Service (AWS SMS) to migrate the remaining servers.

C. Deploy the Migration Evaluator agentless collector to the ESXi hypervisor. Review the collected data in Migration Evaluator. Identify inactive servers. Remove the inactive servers from the migration list. Import the data to AWS Migration Hub.

D. Deploy the AWS Application Migration Service Agent to each VM. When the data is collected, use Amazon Redshift to import and analyze the data. Use Amazon QuickSight for data visualization.

Suggested Answer: *C*

Community vote distribution

kgpoj 10 months, 3 weeks ago

Selected Answer: C

C vs D

Migration Evaluator is suited for initial inventory collection, and it is Agentless so low overhead

In D, the Application Migration Service needs to install agent on thousands of VMs, so it is not suitable for initial inventory collection and is high Operational overhead

upvoted 1 times

igor12ghsj577 1 year, 4 months ago

why to remove highly utilized servers from the list, these answers can be rejected immediately.

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: C

Option C

upvoted 1 times

callmechoice 1 year, 8 months ago

migration evaluator. I think C is correct

upvoted 4 times

NikkyDicky 1 year, 11 months ago

Selected Answer: C

C no doubt

upvoted 1 times

SkyZeroZx 1 year, 12 months ago

Selected Answer: C

C

This solution can meet the requirements with the least operational overhead. and also, keyword for planning only

upvoted 1 times

javitech83 2 years ago

Selected Answer: C

I was first thinking about D because is is stated that the company has little knowledge aboutVMWare. But option D introduces operational overhead

upvoted 1 times

pupsik 2 years ago

Selected Answer: C

C seems like a good choice:

https://aws.amazon.com/migration-evaluator/features/

upvoted 2 times

easytoo 2 years ago

c-c-c-c-c-c-

migration evaluator ftw

upvoted 1 times

easytoo 2 years ago

Question 210 is a-a-a-a-a-a-a-a

upvoted 1 times

yzrk 2 years ago

Selected Answer: C

C

This solution can meet the requirements with the least operational overhead. and also, keyword for planning only

upvoted 3 times

Question #212 *Topic 1*

A company runs a microservice as an AWS Lambda function. The microservice writes data to an on-premises SQL database that supports a limited number of concurrent connections. When the number of Lambda function invocations is too high, the database crashes and causes application downtime. The company has an AWS Direct Connect connection between the company's VPC and the on-premises data center. The company wants to protect the database from crashes.

Which solution will meet these requirements?

A. Write the data to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the Lambda function to read from the queue and write to the existing database. Set a reserved concurrency limit on the Lambda function that is less than the number of connections that the database supports.

B. Create a new Amazon Aurora Serverless DB cluster. Use AWS DataSync to migrate the data from the existing database to Aurora Serverless. Reconfigure the Lambda function to write to Aurora.

C. Create an Amazon RDS Proxy DB instance. Attach the RDS Proxy DB instance to the Amazon RDS DB instance. Reconfigure the Lambda function to write to the RDS Proxy DB instance.

D. Write the data to an Amazon Simple Notification Service (Amazon SNS) topic. Invoke the Lambda function to write to the existing database when the topic receives new messages. Configure provisioned concurrency for the Lambda function to be equal to the number of connections that the database supports.

Suggested Answer: *D*

Community vote distribution

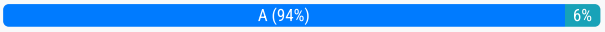

Just_Ninja Highly Voted 1 year, 11 months ago

Selected Answer: A

A tricky question :)
The RDS proxy sounds sexy, but it cannot be used because the database is on premise.

The creative solution here is SQS.
Such questions are partly about your understanding of the services and some solutions are good, even if they sound a bit strange at first :)
upvoted 11 times

joleneinthebackyard Highly Voted 1 year, 8 months ago

Selected Answer: A

"The company wants to protect the database from crashes" means keep the existing one and do something that can prevent crashes, not to migrate it to another in anywhere. -> B, C out

Choice between SQS and SNS is easy.
upvoted 6 times

eesa Most Recent 2 months, 1 week ago

Selected Answer: A

Amazon SQS decouples ingestion from processing, allowing for asynchronous data handling.

By placing data into an SQS queue, you can buffer incoming requests regardless of spikes in Lambda invocation.

You then create a separate Lambda consumer that reads messages from the queue and writes them to the database at a controlled rate.

Using reserved concurrency, you can limit the number of simultaneous Lambda executions to a number lower than the database's connection limit—protecting the database.
upvoted 1 times

bi11 1 year ago

Selected Answer: C

C,

Keyword: "supports a limited number of concurrent connections"

Creating an Amazon RDS Proxy DB instance and attaching it to the Amazon RDS DB instance can help manage the database connections efficiently and prevent the database from being overwhelmed by too many connections. The RDS Proxy can pool and share connections to the database, which can reduce the number of connections that each Lambda function invocation needs to establish. This can help to prevent the database from crashing when the number of Lambda function invocations is high.

Reconfiguring the Lambda function to write to the RDS Proxy DB instance instead of directly to the database can further help to protect the database from crashes. This is because the RDS Proxy can handle the connections to the database, reducing the load on the database and helping to ensure its stability.

upvoted 1 times

altonh 5 months, 1 week ago

You need to first migrate your DB to Amazon RDS, which was never mentioned as one of the steps.

upvoted 1 times

helloworldabc 10 months, 1 week ago

just A

upvoted 1 times

pk0619 1 year, 2 months ago

Selected Answer: A

You can use SQS to write data, however the phrase "reserved concurrency" is incorrect, Lambda has "provisioned concurrency"

upvoted 1 times

TonytheTiger 1 year, 2 months ago

Selected Answer: A

Option A: AWS Tutorial on How To

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/rds-lambda-tutorial.html

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: A

Option A

upvoted 2 times

ggrodskiy 1 year, 11 months ago

Correct A.

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: A

Its an A

upvoted 1 times

javitech83 2 years ago

Selected Answer: A

correct is A as database is on-premises

upvoted 2 times

bhanus 2 years ago

Selected Answer: A

MODERATOR Please delete my previous comment. I commented about RDS proxy which is totally WRONG.

Answer is A

upvoted 1 times

awscerts023 2 years ago

Selected Answer: C

Will go with C , don't think the question says they need to keep the on-prem db

upvoted 1 times

Maria2023 2 years ago

Selected Answer: A

apparently, we need to make the lambda "not to rush that much" and keep the connection within the limit of the on-pre DB. So if we want not to lose data while waiting we implement SQS before the lambda so it keeps the requests in the queue.

upvoted 3 times

SmileyCloud 2 years ago

Selected Answer: A

C should be logical answer, that's what RDS proxy does. But, they want to keep the existing SQL on-prem and not migrate to RDS. So C and B are out. We need to throttle the connections. SNS is not designed for this. So, it's SQS (A).

upvoted 1 times

psyx21 2 years ago

Selected Answer: A

Correct answer is A

upvoted 1 times

easytoo 2 years ago

c-c-c-c-c-c

By creating an Amazon RDS Proxy DB instance and attaching it to the existing Amazon RDS DB instance, you can protect the database from crashes caused by a high number of Lambda function invocations. The RDS Proxy acts as an intermediary between the Lambda function and the database, managing the connections and pooling them efficiently

upvoted 1 times

easytoo 1 year, 11 months ago

a-a-a-a-a-a-a-a

upvoted 3 times

bhanus 2 years ago

Selected Answer: A

A is the answer. RDS proxy is meant to help with connection pooling. Amazon RDS Proxy instance maintains a pool of established connections to your RDS database instances, reducing the stress on database compute and memory resources that typically occurs when new connections are established. RDS Proxy also shares infrequently used database connections, so that fewer connections access the RDS database. This connection pooling enables your database to efficiently support a large number and frequency of application connections so that your application can scale without compromising performance.

upvoted 1 times

bhanus 2 years ago

Answer is A. But IGNORE my above comment on RDS. The current situation is database is on-premises. So RDS proxy has nothing to do with onprem DB. so Answer is A

upvoted 2 times

Question #213 *Topic 1*

A company uses a Grafana data visualization solution that runs on a single Amazon EC2 instance to monitor the health of the company's AWS workloads. The company has invested time and effort to create dashboards that the company wants to preserve. The dashboards need to be highly available and cannot be down for longer than 10 minutes. The company needs to minimize ongoing maintenance.

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

A. Migrate to Amazon CloudWatch dashboards. Recreate the dashboards to match the existing Grafana dashboards. Use automatic dashboards where possible.

B. Create an Amazon Managed Grafana workspace. Configure a new Amazon CloudWatch data source. Export dashboards from the existing Grafana instance. Import the dashboards into the new workspace.

C. Create an AMI that has Grafana pre-installed. Store the existing dashboards in Amazon Elastic File System (Amazon EFS). Create an Auto Scaling group that uses the new AMI. Set the Auto Scaling group's minimum, desired, and maximum number of instances to one. Create an Application Load Balancer that serves at least two Availability Zones.

D. Configure AWS Backup to back up the EC2 instance that runs Grafana once each hour. Restore the EC2 instance from the most recent snapshot in an alternate Availability Zone when required.

Suggested Answer: *B*

Community vote distribution

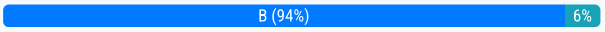

easytoo Highly Voted 2 years ago

Selected Answer: B

By creating an Amazon Managed Grafana workspace, you can offload the operational overhead of managing and maintaining the Grafana infrastructure. Amazon Managed Grafana is a fully managed service that takes care of the underlying infrastructure, including scalability, availability, and updates.

upvoted 6 times

liuliangzhou Most Recent 9 months, 3 weeks ago

Selected Answer: B

The meaning of option B is to create a new Grafana workspace, configure the current CloudWatch data source to it, and then import the historical Grafana instances into the new Grafana workspace.

upvoted 1 times

bacharbhouri 1 year, 1 month ago

Selected Answer: C

The company has invested time and effort to create dashboards that the company wants to preserve.

B is good but it won't preserve their dashboard

upvoted 2 times

bacharbhouri 1 year, 1 month ago

I mean B, sorry. Moderator please change.

upvoted 3 times

surya_lolla 1 year, 6 months ago

Selected Answer: B

Option B is correct, however read this, https://docs.aws.amazon.com/grafana/latest/userguide/AMG-workspace-content-migration.html

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

gotta be a B

upvoted 1 times

SmileyCloud 2 years ago

Selected Answer: B

Def B.

upvoted 1 times

psyx21 2 years ago

Selected Answer: B

Correct answer is B

upvoted 1 times

bhanus 2 years ago

Selected Answer: B

B is the answer https://aws.amazon.com/grafana/

upvoted 3 times

Question #214 *Topic 1*

A company needs to migrate its customer transactions database from on premises to AWS. The database resides on an Oracle DB instance that runs on a Linux server. According to a new security requirement, the company must rotate the database password each year.

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

A. Convert the database to Amazon DynamoDB by using the AWS Schema Conversion Tool (AWS SCT). Store the password in AWS Systems Manager Parameter Store. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly passtard rotation.

B. Migrate the database to Amazon RDS for Oracle. Store the password in AWS Secrets Manager. Turn on automatic rotation. Configure a yearly rotation schedule.

C. Migrate the database to an Amazon EC2 instance. Use AWS Systems Manager Parameter Store to keep and rotate the connection string by using an AWS Lambda function on a yearly schedule.

D. Migrate the database to Amazon Neptune by using the AWS Schema Conversion Tool (AWS SCT). Create an Amazon CloudWatch alarm to invoke an AWS Lambda function for yearly password rotation.

Suggested Answer: *C*

Community vote distribution

B (100%)

joleneinthebackyard Highly Voted 1 year, 2 months ago

Selected Answer: B

Wish all questions are clear like this.
A: Drop immediately at the first sentence
B: sounds good
C: host database in ec2 instance will never a choice. Plus SSM parameter store + lambda for password rotation is not as good as secret manager
D: Again, don't migrate one type of database to another

upvoted 6 times

611c008 Most Recent 7 months, 1 week ago

Selected Answer: B

C is wrong as system manager parm store does not support auto rotate password

upvoted 1 times

kejam 11 months, 2 weeks ago

Selected Answer: B

Answer B
https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database-types-oracle/

upvoted 2 times

career360guru 1 year, 1 month ago

Selected Answer: B

Option B

upvoted 2 times

dkcloudguru 1 year, 3 months ago

Doubt in question it mention yearly rotation, if you can see in Secret Manager the dropdown options are hourly, days, week, and months it doesn't have the yearly option, however, you can mention 12 if that is the case then option B is correct else option C

upvoted 1 times

Simon523 1 year, 4 months ago

Selected Answer: B

https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-other.html#rotate-secrets_turn-on-for-other_step1

upvoted 1 times

Just_Ninja 1 year, 5 months ago

Selected Answer: B

It is sad that so many questions here are marked as correct with a wrong result.

Well Architeting Framework!!!

upvoted 1 times

nicecurls 1 year, 5 months ago

Selected Answer: B

ofc it's B

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

B for sure

upvoted 2 times

Christina666 1 year, 5 months ago

Selected Answer: B

Secrets manager has built-in rotation feature

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: B

keyword = Secrets Manager.

Then B

upvoted 1 times

psyx21 1 year, 6 months ago

Selected Answer: B

Correct answer is B

upvoted 1 times

easytoo 1 year, 6 months ago

b-b-b-b-b-b-b-b

upvoted 1 times

bhanus 1 year, 6 months ago

B is the answer

upvoted 1 times

chiaseed 1 year, 6 months ago

Selected Answer: B

I'd vote for B. A keyword that leads me to B is "rotate the database password each year." This is referring to Secrets Manager.

upvoted 1 times

emiliocb4 1 year, 6 months ago

Selected Answer: B

least operation... rds + secret manager

upvoted 1 times

nexus2020 1 year, 6 months ago

Selected Answer: B

the LEAST operational overhead. So B is the easest

upvoted 2 times

Question #215 *Topic 1*

A solutions architect is designing an AWS account structure for a company that consists of multiple teams. All the teams will work in the same AWS Region. The company needs a VPC that is connected to the on-premises network. The company expects less than 50 Mbps of total traffic to and from the on-premises network.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

A. Create an AWS CloudFormation template that provisions a VPC and the required subnets. Deploy the template to each AWS account.

B. Create an AWS CloudFormation template that provisions a VPC and the required subnets. Deploy the template to a shared services account. Share the subnets by using AWS Resource Access Manager.

C. Use AWS Transit Gateway along with an AWS Site-to-Site VPN for connectivity to the on-premises network. Share the transit gateway by using AWS Resource Access Manager.

D. Use AWS Site-to-Site VPN for connectivity to the on-premises network.

E. Use AWS Direct Connect for connectivity to the on-premises network.

Suggested Answer: *AD*

Community vote distribution

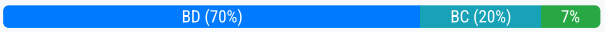

NikkyDicky Highly Voted 1 year, 11 months ago

Selected Answer: BD

BD
they need a (one) VPC, no need for TGW.
Use case for subnet sharing via RAM
upvoted 13 times

LuongTo 7 months ago
why A out?
upvoted 1 times

KennethYY 6 months ago
because deploy to "each account"
upvoted 1 times

8693a49 Most Recent 11 months ago

Selected Answer: AC

They are designing an account structure. This means multiple accounts, implicitly multiple VPCs. So A will take care of account provisioning. (B is incorrect, subnets cannot be shared). To connect to on-prem, site-to-site VPN is sufficient and most cost-effective, and we also need to give access to it from all accounts, so we need a Transit Gateway. Therefore C is the other correct answer. (D is incorrect because it only works for one VPC, one account, and E is incorrect because is more expensive than VPN and not necessary)
upvoted 2 times

helloworldabc 10 months, 1 week ago
just BD
upvoted 1 times

8693a49 11 months ago
Correction. VPC subnets can be shared, so BC would work, but the resulting architecture is a networking nightmare. I would not do that.
upvoted 2 times

0dc6cac 2 weeks ago
The question says "cost-effective", transit gateways + multiple site-to-site VPNs are not cheap. In most cases, B should be enough, so you only need one site-to-site connection
upvoted 1 times

helloworldabc 10 months, 1 week ago
Transit gateways are not cost-effective
upvoted 1 times

gfhbox0083 11 months, 3 weeks ago

B, D for sure.

No need for a TGW

upvoted 1 times

LuongTo 7 months ago

why A out?

upvoted 1 times

bacharbhouri 1 year, 1 month ago

Selected Answer: BE

Why is nobody considering Direct Connect, it is cheaper than Site to Site VPN.

upvoted 1 times

bacharbhouri 1 year, 1 month ago

the ask here is for most cost effectively choice.

upvoted 1 times

YOUSSEFSWAID 1 year, 1 month ago

If you have one VPC why you need to share the subnets ?

upvoted 2 times

TonytheTiger 1 year, 3 months ago

Selected Answer: BD

Option BC & NOT C - The MOST cost effective option: AWS Site-to-Site VPN connection pricing still applies in addition to AWS Transit Gateway VPN attachment pricing. So you will be additional cost with both option

https://aws.amazon.com/transit-gateway/pricing/

upvoted 2 times

ftaws 1 year, 5 months ago

The problem did not say how many VPC. @@@

upvoted 2 times

pk0619 6 months, 1 week ago

there is just one VPC if you select B which makes D the right choice for second answer

upvoted 1 times

ayadmawla 1 year, 6 months ago

Selected Answer: BC

B+C in my humble opinion. Reason for C is that this is a design for a company with "multiple teams" so it is only logical that these teams will want to have at some stage independent accounts from one another and different accounts within the same teams. Thinking about a single VPC would be a bit short sighted.

upvoted 3 times

career360guru 1 year, 7 months ago

Selected Answer: BD

B and D is right choice.

upvoted 2 times

lghoshino78 1 year, 7 months ago

Selected Answer: AD

Most Cost Effective...

upvoted 1 times

nublit 1 year, 7 months ago

Selected Answer: AD

You need to create a singe VPC and a single Account.

upvoted 1 times

SK_Tyagi 1 year, 10 months ago

Selected Answer: BD

Direct Connect may be an overkill with 1GBPs

upvoted 3 times

kebmiockey 1 year, 10 months ago

Other problem with VPN is 1.25 Gb limitation.

upvoted 1 times

ggrodskiy 1 year, 11 months ago

Correct AD.

I think A is correct because you can connect the VPN to each VPC by using a VPN connection resource in each AWS account. You do not need a shared network account for that. You can refer to this documentation for more details:

https://docs.aws.amazon.com/vpn/latest/s2svpn/VPC_VPN.html

B is not correct because it will create a single VPC for all the AWS accounts, which will reduce the isolation and security for the different teams. It will also require sharing the subnets by using AWS Resource Access Manager, which will add complexity and overhead.

upvoted 3 times

Christina666 1 year, 12 months ago

Selected Answer: BD

Tgw is for VPCs communication.

upvoted 1 times

SmileyCloud 1 year, 12 months ago

Selected Answer: BC

BC. There are multiple teams and accounts.

upvoted 3 times

SkyZeroZx 1 year, 12 months ago

Selected Answer: BD

BD? dont think we need tgw here.

upvoted 1 times

Question #216 *Topic 1*

A solutions architect at a large company needs to set up network security for outbound traffic to the internet from all AWS accounts within an organization in AWS Organizations. The organization has more than 100 AWS accounts, and the accounts route to each other by using a centralized AWS Transit Gateway. Each account has both an internet gateway and a NAT gateway for outbound traffic to the internet. The company deploys resources only into a single AWS Region.

The company needs the ability to add centrally managed rule-based filtering on all outbound traffic to the internet for all AWS accounts in the organization. The peak load of outbound traffic will not exceed 25 Gbps in each Availability Zone.

Which solution meets these requirements?

A. Create a new VPC for outbound traffic to the internet. Connect the existing transit gateway to the new VPC. Configure a new NAT gateway. Create an Auto Scaling group of Amazon EC2 instances that run an open-source internet proxy for rule-based filtering across all Availability Zones in the Region. Modify all default routes to point to the proxy's Auto Scaling group.

B. Create a new VPC for outbound traffic to the internet. Connect the existing transit gateway to the new VPC. Configure a new NAT gateway. Use an AWS Network Firewall firewall for rule-based filtering. Create Network Firewall endpoints in each Availability Zone. Modify all default routes to point to the Network Firewall endpoints.

C. Create an AWS Network Firewall firewall for rule-based filtering in each AWS account. Modify all default routes to point to the Network Firewall firewalls in each account.

D. In each AWS account, create an Auto Scaling group of network-optimized Amazon EC2 instances that run an open-source internet proxy for rule-based filtering. Modify all default routes to point to the proxy's Auto Scaling group.

Suggested Answer: *D*

Community vote distribution

B (100%)

bjexamprep Highly Voted 9 months, 2 weeks ago

Selected Answer: B

Centrally managed egress, so C/D are out.
Both A and B are wrong, because
1. There isn't internet gateway.
2. "Modify all default routes to point to the …". A firewall or "proxy's Auto Scaling group" don't have public IP, the default route must be pointing to the NAT gateway. And NAT gateway has a peer public IP configured on the IGW. The route should be: internet prefix of all the internal subnet-> NAT gateway -> firewall -> internet gateway, and reverse routing rules are also required.
.
Well, considering the persistent low quality of AWS Exam Questions, I vote B

upvoted 5 times

easytoo Highly Voted 1 year, 6 months ago

b-b-b-b-b-b

Create a new VPC specifically dedicated to outbound traffic to the internet. This helps isolate and manage the outbound traffic separately from other resources.
Connect the existing transit gateway to the new VPC. This ensures that the VPC is connected to the centralized transit gateway that routes traffic between AWS accounts.
Configure a new NAT gateway within the new VPC. This NAT gateway provides the necessary outbound connectivity to the internet for resources within the VPC.
Use AWS Network Firewall, a managed firewall service, for rule-based filtering on the outbound traffic. Network Firewall allows you to define and enforce custom rules for traffic leaving the VPC.
Create Network Firewall endpoints in each Availability Zone. These endpoints serve as the traffic inspection points where Network Firewall applies the filtering rules.
Modify all default routes in the VPCs to point to the Network Firewall endpoints. This ensures that all outbound traffic from the VPCs flows through the Network Firewall for rule-based filtering.

upvoted 5 times

thotwielder Most Recent 9 months, 3 weeks ago

Selected Answer: B

c,d in each AWS account. wrong

a: use third party solution, not as good as b (use aws service)

upvoted 2 times

career360guru 1 year, 1 month ago

Selected Answer: B

Option B

upvoted 1 times

rlf 1 year, 2 months ago

B.

https://docs.aws.amazon.com/whitepapers/latest/building-scalable-secure-multi-vpc-network-infrastructure/using-nat-gateway-with-firewall.html

upvoted 4 times

duriselvan 1 year, 3 months ago

https://aws.amazon.com/blogs/security/hands-on-walkthrough-of-the-aws-network-firewall-flexible-rules-engine/

upvoted 2 times

xav1er 1 year, 4 months ago

Selected Answer: B

Given the available options and the requirements:

B. Create an interface VPC endpoint for API Gateway, and set an endpoint policy to only allow access to the specific API. Add a resource policy to API Gateway to only allow access from the VPC endpoint. Change the API Gateway endpoint type to private. is the correct answer.

upvoted 1 times

chikorita 1 year, 4 months ago

bro what?

upvoted 2 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

B for sure

upvoted 1 times

Christina666 1 year, 5 months ago

Selected Answer: B

centrally managed outbound traffic: tgw-> centralized VPC with network firewall with rules-> internet

upvoted 4 times

chiaseed 1 year, 6 months ago

Selected Answer: B

vote for B. The keyword is "centrally managed rule-based filtering on outbound traffic to the internet for all AWS accounts...". Network Firewall can centrally manage network security policies.

upvoted 3 times

SmileyCloud 1 year, 6 months ago

Selected Answer: B

B. Answer A is similar, but you have to deal with EC2 instances and dealing with 3rd party FW, not good - management overhead. C is impossible. D is waay to much hard to manage.

upvoted 2 times

psyx21 1 year, 6 months ago

Selected Answer: B

Correct answer is B

upvoted 1 times

nexus2020 1 year, 6 months ago

Selected Answer: B

vote for B

upvoted 2 times

Question #217 *Topic 1*

A company uses a load balancer to distribute traffic to Amazon EC2 instances in a single Availability Zone. The company is concerned about security and wants a solutions architect to re-architect the solution to meet the following requirements:

• Inbound requests must be filtered for common vulnerability attacks.
• Rejected requests must be sent to a third-party auditing application.
• All resources should be highly available.

Which solution meets these requirements?

A. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Use Amazon Inspector to monitor traffic to the ALB and EC2 instances. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB. Use an AWS Lambda function to frequently push the Amazon Inspector report to the third-party auditing application.

B. Configure an Application Load Balancer (ALB) and add the EC2 instances as targets. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB name and enable logging with Amazon CloudWatch Logs. Use an AWS Lambda function to frequently push the logs to the third-party auditing application.

C. Configure an Application Load Balancer (ALB) along with a target group adding the EC2 instances as targets. Create an Amazon Kinesis Data Firehose with the destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the web ACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

D. Configure a Multi-AZ Auto Scaling group using the application's AMI. Create an Application Load Balancer (ALB) and select the previously created Auto Scaling group as the target. Create an Amazon Kinesis Data Firehose with a destination of the third-party auditing application. Create a web ACL in WAF. Create an AWS WAF using the WebACL and ALB then enable logging by selecting the Kinesis Data Firehose as the destination. Subscribe to AWS Managed Rules in AWS Marketplace, choosing the WAF as the subscriber.

Suggested Answer: *B*

Community vote distribution

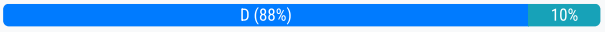

Maria2023 Highly Voted 2 years ago

Selected Answer: D

Only A and D cover the requirement for high availability. A uses Inspector, which is a vulnerability scanner and does not monitor traffic. So - even that I don't like the complexity of D - this remains the only option

upvoted 16 times

SK_Tyagi Highly Voted 1 year, 10 months ago

Selected Answer: D

I was confused between A and D, but seems WAF can deliver logs to Firehose
https://docs.aws.amazon.com/waf/latest/developerguide/logging-kinesis.html

upvoted 6 times

85b5b55 Most Recent 4 weeks, 1 day ago

Selected Answer: D

A and D look right. But D is the correct answer. As A is using Amazon Inspector, it doesn't support monitoring traffic flow.

upvoted 1 times

liuliangzhou 9 months, 3 weeks ago

Selected Answer: D

Compared to A, prioritize AWS Kinesis over third-party auditing applications

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: B

D is good option but as the question does not mention about 3rd party auditing app it may not be possible to directly integrate it with Firehose. One may have to use http api to push the logs - as this is not mentioned I will go with Option B.

upvoted 1 times

career360guru 1 year, 7 months ago

Oh Mistake, I want to change it to D as B does not support High Availability.

upvoted 2 times

xav1er 1 year, 10 months ago

Selected Answer: D

It's D, makes most sense,

upvoted 2 times

chico2023 1 year, 10 months ago

This is such a mal formed question...

You see, nowhere in the question we are told about customer's application. However we are told they want ALL their resources highly available. B would be sooo much better if there wasn't that "All resources should be highly available." because, seriously, D is not the best in my opinion. We don't know much what applications they use, what third party auditing application and so on...

Anyway, it might be D after all, but oh my...

upvoted 2 times

ggrodskiy 1 year, 11 months ago

Correct D.

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: D

its a D

upvoted 1 times

javitech83 2 years ago

Selected Answer: D

ASG in Multiple AZ. WAF and WAF logs with kinesis

upvoted 1 times

chikorita 2 years ago

"enable logging by selecting the Kinesis Data Firehose as the destination"--- how can ALB write logs directly to Kinesis???

it should be CW logs group

any links for help??

upvoted 1 times

Masonyeoh 2 years ago

Selected Answer: D

Amazon inspector does NOT inspect traffic coming to an Application Load Balancer (ALB)

upvoted 3 times

PhuocT 2 years ago

Selected Answer: D

D is correct answer

Inbound requests must be filtered for common vulnerability attacks -> WAF

Rejected requests must be sent to a third-party auditing application-> Enable access log and use kinesis stream to send logs to third party

All resources should be highly available -> Muti AZ auto scaling group.

upvoted 4 times

ozelllll 2 years ago

Selected Answer: D

Inspector does not filter inbound traffic for attack signatures, this is what WAF is for

upvoted 2 times

SmileyCloud 2 years ago

Selected Answer: A

B and C do not provide HA. D is similar to A but lacks Inspector -> "Amazon Inspector automatically discovers workloads, such as Amazon EC2 instances, containers, and Lambda functions, and scans them for software vulnerabilities and unintended network exposure."

upvoted 2 times

javitech83 2 years ago

but you need logs of the reject request on WAF. So I think correct answer is D
upvoted 1 times

SmileyCloud 1 year, 11 months ago
It's probably B. C and D are not correct, ALB can't send logs to Kinesis Fire Hose.
upvoted 1 times

easytoo 2 years ago
a-a-a-a-a-a-a-a
multi-az for HA
upvoted 1 times

easytoo 1 year, 11 months ago
it's d-d-d-d-d-d--d-d
upvoted 1 times

bhanus 2 years ago
Selected Answer: D
I got with D. The reason to go with D is because other options ABC are wrong.
1. It says use Amazon Inspector to inspect traffic to ALB. This is wrong. Amazon inspector does NOT inspect traffic coming to an Application Load Balancer (ALB). Amazon Inspector is a security assessment service that helps you analyze the security and compliance of your EC2 instances and applications running on them. To inspect traffic coming to an ALB, you can consider using other services such as AWS WAF (Web Application Firewall) or AWS Shield. AWS WAF allows you to define rules to filter and block malicious traffic targeting your ALB.
B - Does NOT talk about HA as it is asked in ques
C - Does NOT talk about HA as it is asked in ques
upvoted 3 times

Question #218 *Topic 1*

A company is running an application in the AWS Cloud. The application consists of microservices that run on a fleet of Amazon EC2 instances in multiple Availability Zones behind an Application Load Balancer. The company recently added a new REST API that was implemented in Amazon API Gateway. Some of the older microservices that run on EC2 instances need to call this new API.

The company does not want the API to be accessible from the public internet and does not want proprietary data to traverse the public internet.

What should a solutions architect do to meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the VPC and the API Gateway. Use API Gateway to generate a unique API Key for each microservice. Configure the API methods to require the key.

B. Create an interface VPC endpoint for API Gateway, and set an endpoint policy to only allow access to the specific API. Add a resource policy to API Gateway to only allow access from the VPC endpoint. Change the API Gateway endpoint type to private.

C. Modify the API Gateway to use IAM authentication. Update the IAM policy for the IAM role that is assigned to the EC2 instances to allow access to the API Gateway. Move the API Gateway into a new VPDeploy a transit gateway and connect the VPCs.

D. Create an accelerator in AWS Global Accelerator, and connect the accelerator to the API Gateway. Update the route table for all VPC subnets with a route to the created Global Accelerator endpoint IP address. Add an API key for each service to use for authentication.

Suggested Answer: *C*

Community vote distribution

B (100%)

SkyZeroZx Highly Voted 12 months ago

Selected Answer: B

Tip: Anytime you see "don't want to traverse Internet traffic" always look for endpoint in the answers. Most likely, that's the answer.

upvoted 11 times

Just_Ninja Highly Voted 11 months, 1 week ago

Selected Answer: B

The quality control here is unfortunately not as expected when you buy access.
C is due nonsense.
B is correct.
VPC Endpoint to API Gateway and a policy on both sides!

Trust me, i´m a Ninja

upvoted 6 times

rxhan 11 months, 1 week ago

thanks Ninja

upvoted 2 times

shaaam80 Most Recent 7 months ago

Selected Answer: B

Answer B - VPC Interface endpoint to privately access services without data over internet.

upvoted 3 times

career360guru 7 months, 1 week ago

Selected Answer: B

Option B

upvoted 1 times

NikkyDicky 11 months, 4 weeks ago

Selected Answer: B

B for sure

upvoted 1 times

Alabi 1 year ago

Selected Answer: B

B for sure

upvoted 1 times

SmileyCloud 1 year ago

Selected Answer: B

Tip: Anytime you see "don't want to traverse Internet traffic" always look for endpoint in the answers. Most likely, that's the answer.

upvoted 3 times

easytoo 1 year ago

b-b-b-b-b-b-b

By implementing this solution, the company can ensure that the new API in API Gateway is not accessible from the public internet. The interface VPC endpoint provides private connectivity, allowing secure communication between the microservices running on EC2 instances and the API Gateway. This ensures the proprietary data does not traverse the public internet, enhancing security and data protection.

upvoted 3 times

bhanus 1 year ago

I vote B

upvoted 1 times

nexus2020 1 year ago

Selected Answer: B

VPC endpoint usualy is the prefect answer to avoid internet traffic

upvoted 1 times

Question #219 *Topic 1*

A company has set up its entire infrastructure on AWS. The company uses Amazon EC2 instances to host its ecommerce website and uses Amazon S3 to store static data. Three engineers at the company handle the cloud administration and development through one AWS account. Occasionally, an engineer alters an EC2 security group configuration of another engineer and causes noncompliance issues in the environment.

A solutions architect must set up a system that tracks changes that the engineers make. The system must send alerts when the engineers make noncompliant changes to the security settings for the EC2 instances.

What is the FASTEST way for the solutions architect to meet these requirements?

A. Set up AWS Organizations for the company. Apply SCPs to govern and track noncompliant security group changes that are made to the AWS account.

B. Enable AWS CloudTrail to capture the changes to EC2 security groups. Enable Amazon CloudWatch rules to provide alerts when noncompliant security settings are detected.

C. Enable SCPs on the AWS account to provide alerts when noncompliant security group changes are made to the environment.

D. Enable AWS Config on the EC2 security groups to track any noncompliant changes. Send the changes as alerts through an Amazon Simple Notification Service (Amazon SNS) topic.

Suggested Answer: *B*

Community vote distribution

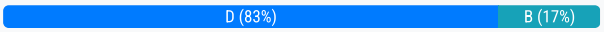

Soweetadad Highly Voted 1 year, 10 months ago

Selected Answer: D

Both B and D work, except B has no notification set.
https://aws.amazon.com/blogs/security/how-to-monitor-aws-account-configuration-changes-and-api-calls-to-amazon-ec2-security-groups/
upvoted 10 times

bhanus Highly Voted 2 years ago

Selected Answer: D

I vote D. aws config changes can be sent to SNS topic https://docs.aws.amazon.com/config/latest/developerguide/notifications-for-AWS-Config.html
upvoted 6 times

ry1999 Most Recent 9 months, 3 weeks ago

Selected Answer: B

B is faster
upvoted 1 times

kgpoj 10 months, 3 weeks ago

Selected Answer: D

Both B and D works. But the question is asking for FASTEST.

For cloudTrail, you need: CloudTrail → CloudWatch Logs → CloudWatch Metric Filter → CloudWatch Alarm → SNS Notification

For aws Config, it natively support integration with SNS.
Hence we should choose D
upvoted 3 times

skipbaylessfor3 11 months ago

I'm leaning towards D, but looks what it says in this blog:
https://aws.amazon.com/blogs/security/how-to-monitor-aws-account-configuration-changes-and-api-calls-to-amazon-ec2-security-groups/

For the Config option, it says:
"The use of AWS Config in Method 1 allows for the configuration of a security group to be tracked along with other AWS resources. Changes to the security group's configuration are reported during the next Config compliance evaluation, typically within 10 minutes"

and for the CloudTrail option it says:
"The use of CloudTrail and CloudWatch Events in Method 2 allows for the near real-time detection of API calls that could change the configuration of a VPC security group"

So it seems clear cut to me that the answer is B, although if I hadn't seen this blog I would've picked D probably
upvoted 1 times

red_panda 1 year, 1 month ago

Selected Answer: B

For me the answer is B.
Here we are talking about "tracking al changes" and "notify for non-compliant".
It's certainly a very ambiguous question that the folks at AWS could have spared us, but for me (and for chat-gpt) B is the answer :)
upvoted 3 times

9esh 1 year, 3 months ago

D: AWS Config provides rules to detect non-complaint config
B: Can track all event however doesn't provide native support for rules to detect non-complaint changes
upvoted 1 times

dankositzke 1 year, 4 months ago

Selected Answer: B

In my opinion, the question asks for (1) a “system that tracks CHANGES” and (2) asks to “send alerts when the engineers make NONCOMPLIANT CHANGES,” I would choose B since B satisfies the first condition and D does not.

B: implies that CloudTrail tracks all changes.
D: states that Config will only track noncompliant changes, but question is asking for all changes.

But overall this is just another poorly constructed and ambiguous question and answer, which seems to be the norm with these lol
upvoted 1 times

helloworldabc 10 months, 1 week ago

just D
upvoted 1 times

fartosh 1 year, 1 month ago

Actually, AWS Config cannot track *only* non-compliant changes, it always tracks all changes against monitored resources - that's by design. You set rules in AWS Config that indicate whether the change is compliant, but all the changes must be recorded.
https://docs.aws.amazon.com/config/latest/developerguide/how-does-config-work.html#resource-tracking
upvoted 1 times

duriselvan 1 year, 6 months ago

B is ans
https://aws.amazon.com/blogs/security/how-to-monitor-aws-account-configuration-changes-and-api-calls-to-amazon-ec2-security-groups/
Speed: Implementing CloudTrail and CloudWatch is faster than setting up AWS Organizations or using SCPs. You can do it in minutes without modifying the entire account structure or deploying additional resources.
Granularity: CloudTrail and CloudWatch offer fine-grained control over monitoring and alerting, allowing you to define specific rules for noncompliant security settings.
Flexibility: You can easily adapt the CloudWatch rules to different types of noncompliance and adjust the alerts to suit your notification needs.
Existing infrastructure: If the company already uses CloudTrail for logging, setting up CloudWatch rules is a natural extension without requiring significant changes.
upvoted 2 times

shaaam80 1 year, 7 months ago

Selected Answer: D

Answer D. AWS Config is perfect to track config changes. SNS for notification.
upvoted 4 times

career360guru 1 year, 7 months ago

Selected Answer: B

B is better option than D. D only sends an SNS alert when there are non-compliant changes. It does not allow you to actually track each and every changes engineers make.
upvoted 2 times

Jay_2pt0_1 1 year, 7 months ago

I thought so too, initially, but as others have said, B does not actually send the alert.
upvoted 2 times

ghadxx 1 year, 10 months ago
It's D
https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html
upvoted 2 times

ggrodskiy 1 year, 11 months ago
Correct D.
upvoted 1 times

NikkyDicky 1 year, 11 months ago
Selected Answer: D
D works and faster
B would work with adding a CW alert, but D still better
upvoted 4 times

javitech83 2 years ago
Selected Answer: D
correct is D
upvoted 2 times

SkyZeroZx 2 years ago
Selected Answer: D
D
reference link
https://aws.amazon.com/es/blogs/industries/how-to-monitor-alert-and-remediate-non-compliant-hipaa-findings-on-aws/
upvoted 5 times

SmileyCloud 2 years ago
Selected Answer: D
It's D. Check this link, something similar: https://aws.amazon.com/blogs/industries/how-to-monitor-alert-and-remediate-non-compliant-hipaa-findings-on-aws/
upvoted 4 times

Question #220 *Topic 1*

A company has IoT sensors that monitor traffic patterns throughout a large city. The company wants to read and collect data from the sensors and perform aggregations on the data.

A solutions architect designs a solution in which the IoT devices are streaming to Amazon Kinesis Data Streams. Several applications are reading from the stream. However, several consumers are experiencing throttling and are periodically encountering a ReadProvisionedThroughputExceeded error.

Which actions should the solutions architect take to resolve this issue? (Choose three.)

A. Reshard the stream to increase the number of shards in the stream.

B. Use the Kinesis Producer Library (KPL). Adjust the polling frequency.

C. Use consumers with the enhanced fan-out feature.

D. Reshard the stream to reduce the number of shards in the stream.

E. Use an error retry and exponential backoff mechanism in the consumer logic.

F. Configure the stream to use dynamic partitioning.

Suggested Answer: *ACE*

Community vote distribution

ACE (100%)

easytoo Highly Voted 1 year, 6 months ago

To resolve the issue of throttling and ReadProvisionedThroughputExceeded errors in the Amazon Kinesis Data Streams scenario, the solutions architect should take the following actions:

1. A. Reshard the stream to increase the number of shards in the stream: By increasing the number of shards, you can increase the overall throughput capacity of the stream, allowing for more concurrent consumers to read from the stream without being throttled.

2. C. Use consumers with the enhanced fan-out feature: Enhanced fan-out allows for multiple consumers to read from the same shard concurrently, without being limited by the read capacity of the shard. This helps distribute the load and reduces the chances of throttling.

3. E. Use an error retry and exponential backoff mechanism in the consumer logic: Implementing an error retry mechanism with exponential backoff in the consumer logic will help handle throttling errors gracefully. When a ReadProvisionedThroughputExceeded error occurs, the consumer can retry the read operation after a certain delay, gradually increasing the delay between retries to avoid overwhelming the system.

upvoted 19 times

yorkicurke Highly Voted 1 year, 2 months ago

Selected Answer: ACE

this link will explain it all. looks like this question was taken from here.
https://repost.aws/knowledge-center/kinesis-readprovisionedthroughputexceeded

upvoted 9 times

shaaam80 Most Recent 1 year, 1 month ago

Selected Answer: ACE

Answer ACE

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: ACE

A, C, E Options

upvoted 1 times

totten 1 year, 2 months ago

Selected Answer: ACE

Option (D) "Reshard the stream to reduce the number of shards" is generally not a recommended solution because it reduces the capacity of the stream, which might lead to more throttling issues. Reducing shards should only be considered if you're overprovisioned, and reducing capacity will

not negatively impact your consumers.

Option (B) "Use the Kinesis Producer Library (KPL) and adjust the polling frequency" may not be directly related to solving the throttling issue. The KPL is primarily used for producing data into the Kinesis stream, not consuming it.

Option (F) "Configure the stream to use dynamic partitioning" can be beneficial for even distribution of data but is not directly related to resolving throttling issues. Dynamic partitioning is more about balancing the data across shards and does not increase overall read capacity.

So, the most relevant actions to address the throttling issue are (A), (C), and (E).
upvoted 5 times

GoKhe 1 year ago
Nice way to explain the reasons other way round :-)
upvoted 1 times

ggrodskiy 1 year, 5 months ago
Correct ACE.
upvoted 1 times

NikkyDicky 1 year, 5 months ago
Selected Answer: ACE
ACE it
upvoted 1 times

SkyZeroZx 1 year, 5 months ago
Selected Answer: ACE
ACE is correct
upvoted 1 times

SmileyCloud 1 year, 6 months ago
Selected Answer: ACE
Eliminate B, KPL is for writing. "The Kinesis Producer Library (KPL) simplifies producer application development, allowing developers to achieve high write throughput to a Kinesis data stream. " The error was reading.

F, dynamic partitioning is used for different use cases.https://docs.aws.amazon.com/firehose/latest/dev/dynamic-partitioning.html
upvoted 3 times

psyx21 1 year, 6 months ago
Selected Answer: ACE
ACE is correct
upvoted 1 times

nexus2020 1 year, 6 months ago
Selected Answer: ACE
not sure about E, but I would go with AC
upvoted 1 times

Question #221 *Topic 1*

A company uses AWS Organizations to manage its AWS accounts. The company needs a list of all its Amazon EC2 instances that have underutilized CPU or memory usage. The company also needs recommendations for how to downsize these underutilized instances.

Which solution will meet these requirements with the LEAST effort?

A. Install a CPU and memory monitoring tool from AWS Marketplace on all the EC2 instances. Store the findings in Amazon S3. Implement a Python script to identify underutilized instances. Reference EC2 instance pricing information for recommendations about downsizing options.

B. Install the Amazon CloudWatch agent on all the EC2 instances by using AWS Systems Manager. Retrieve the resource optimization recommendations from AWS Cost Explorer in the organization's management account. Use the recommendations to downsize underutilized instances in all accounts of the organization.

C. Install the Amazon CloudWatch agent on all the EC2 instances by using AWS Systems Manager. Retrieve the resource optimization recommendations from AWS Cost Explorer in each account of the organization. Use the recommendations to downsize underutilized instances in all accounts of the organization.

D. Install the Amazon CloudWatch agent on all the EC2 instances by using AWS Systems Manager. Create an AWS Lambda function to extract CPU and memory usage from all the EC2 instances. Store the findings as files in Amazon S3. Use Amazon Athena to find underutilized instances. Reference EC2 instance pricing information for recommendations about downsizing options.

Suggested Answer: *B*

Community vote distribution

Maria2023 Highly Voted 1 year, 6 months ago

Actually, the right answer is to use Compute Optimizer, I don't understand why it was not part of the choices here
https://aws.amazon.com/compute-optimizer/

upvoted 13 times

totten Highly Voted 1 year, 2 months ago

Selected Answer: B

AWS Cost Explorer provides resource optimization recommendations, including rightsizing EC2 instances based on historical usage data. These recommendations are generated for each account in the organization's management account, so you can obtain insights for all accounts centrally.

Option A introduces complexity by requiring the company to install a third-party tool on all EC2 instances, and then manually develop and maintain a custom script for identifying underutilized instances.

Option C would require you to retrieve recommendations separately for each account within the organization, increasing the administrative overhead compared to a centralized management approach.

Option D, while using native AWS services for data collection, involves creating and maintaining additional AWS services, which is more complex than the straightforward combination of CloudWatch and AWS Cost Explorer.

upvoted 7 times

duriselvan Most Recent 1 year ago

Let's analyze each option based on effort:

A. Marketplace tool:

Effort: High
Requires manual installation of a third-party tool on all instances.
Needs custom script development to identify underutilized instances.
Manual effort needed to reference pricing information for downsizing.
B. Cost Explorer in Org Management Account:

Effort: Low
Leverages existing tools (CloudWatch agent & Cost Explorer) already available.

Recommendations readily available in the management account.
Downsizing options directly available within Cost Explorer.
upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: B

Option B
upvoted 1 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: B

IMO it could be done with either B or D. But the differentiator is "Least Effort" that makes it B
upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

its a B
upvoted 1 times

bhanus 1 year, 6 months ago

Selected Answer: B

Though I vote B. No better choice. This is worst ques. How can cost explorer provide recommendations?. Its should be cost optimizer
upvoted 3 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: B

Classic usage de Cloudwatch metrics and AWS Organization in master account .
C not because more overhead each account for example 100 accounts.
Note : Compute Optimizer is more apropiate in this case but no exist option
upvoted 1 times

easytoo 1 year, 6 months ago

B. Install the Amazon CloudWatch agent on all the EC2 instances using AWS Systems Manager. Retrieve the resource optimization recommendations from AWS Cost Explorer in the organization's management account. Use the recommendations to downsize underutilized instances in all accounts of the organization.

This solution leverages the capabilities of AWS CloudWatch and AWS Cost Explorer to monitor and analyze the CPU and memory usage of EC2 instances. By installing the CloudWatch agent, you can collect the necessary metrics for monitoring. AWS Cost Explorer provides resource optimization recommendations, which can be accessed from the organization's management account. These recommendations can then be used to identify underutilized instances and make informed decisions about downsizing.

This solution requires minimal effort as it utilizes existing AWS services and tools, eliminating the need for additional installations or custom scripts. It also provides a centralized approach by retrieving recommendations from the organization's management account, allowing for efficient management of all accounts within the organization.
upvoted 2 times

SmileyCloud 1 year, 6 months ago

Selected Answer: B

B. That's why you have the management account so you don't have to go to 1000+ accounts and get metrics.
upvoted 3 times

bhanus 1 year, 6 months ago

Selected Answer: B

B - Management account is the key word
upvoted 1 times

nexus2020 1 year, 6 months ago

Selected Answer: B

B. the standard way AWS recommended
upvoted 1 times

Question #222 *Topic 1*

A company wants to run a custom network analysis software package to inspect traffic as traffic leaves and enters a VPC. The company has deployed the solution by using AWS CloudFormation on three Amazon EC2 instances in an Auto Scaling group. All network routing has been established to direct traffic to the EC2 instances.

Whenever the analysis software stops working, the Auto Scaling group replaces an instance. The network routes are not updated when the instance replacement occurs.

Which combination of steps will resolve this issue? (Choose three.)

A. Create alarms based on EC2 status check metrics that will cause the Auto Scaling group to replace the failed instance.

B. Update the CloudFormation template to install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to send process metrics for the application.

C. Update the CloudFormation template to install AWS Systems Manager Agent on the EC2 instances. Configure Systems Manager Agent to send process metrics for the application.

D. Create an alarm for the custom metric in Amazon CloudWatch for the failure scenarios. Configure the alarm to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

E. Create an AWS Lambda function that responds to the Amazon Simple Notification Service (Amazon SNS) message to take the instance out of service. Update the network routes to point to the replacement instance.

F. In the CloudFormation template, write a condition that updates the network routes when a replacement instance is launched.

Suggested Answer: *ADF*

Community vote distribution

NikkyDicky Highly Voted 1 year, 11 months ago

Selected Answer: BDE

CW agent->CW metric->CW alarm->Lambda action

upvoted 9 times

bjexamprep Highly Voted 1 year, 4 months ago

Selected Answer: BDE

This is a bad question design.

The question is looking for a solution for “The network routes are not updated when the instance replacement occurs.”, which means the ASG already has the capability to detect the failure node. With this assumption, there is NO need to install a CloudWatch agent on the EC2 instance, cause the CloudWatch agent in B is doing the same thing.

The correct solution is to use the ASG Lifecycle Hook to invoke the Lambda to update the route.

A better solution is to create a loadbalancer targeting the ASG, and update the route to point to the loadbalancer. With this solution, there is no need to update the route anymore.

upvoted 8 times

chris_spencer Most Recent 8 months, 2 weeks ago

Selected Answer: BDE

BDE.. but a professional should use ASG Lifecycle hooks https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html

upvoted 1 times

NoDoubkevo 9 months, 2 weeks ago

you cannot update templates you can version them.

ADE

upvoted 1 times

chris_spencer 8 months, 2 weeks ago

why can't you update CloudFormation templates?

upvoted 1 times

shaaam80 1 year, 7 months ago

Answer - BDE
Install CW agent on all instances using CF template
Configure CW to send out metrics to SNS
Configure Lambda as SNS target to terminate instance and update n/w routes on the new instances
upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: BDE

B, D, E
upvoted 2 times

Piccaso 1 year, 12 months ago

Selected Answer: BDE

A and F must be wrong.
upvoted 2 times

PhuocT 2 years ago

Selected Answer: BDE

B, D and E
upvoted 3 times

easytoo 2 years ago

b-d-e seems reasonable.
upvoted 2 times

SmileyCloud 2 years ago

Selected Answer: BDE

A is redundant because "Whenever the analysis software stops working, the Auto Scaling group replaces an instance."
C is not correct. AWS System Manager Agebt is not used "to send process metrics for the application."

So, B, D and E because they make a flow.
upvoted 4 times

james55 2 years ago

Selected Answer: BDE

b----d----e
upvoted 1 times

Question #223 *Topic 1*

A company is developing a new on-demand video application that is based on microservices. The application will have 5 million users at launch and will have 30 million users after 6 months. The company has deployed the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. The company developed the application by using ECS services that use the HTTPS protocol.

A solutions architect needs to implement updates to the application by using blue/green deployments. The solution must distribute traffic to each ECS service through a load balancer. The application must automatically adjust the number of tasks in response to an Amazon CloudWatch alarm.

Which solution will meet these requirements?

A. Configure the ECS services to use the blue/green deployment type and a Network Load Balancer. Request increases to the service quota for tasks per service to meet the demand.

B. Configure the ECS services to use the blue/green deployment type and a Network Load Balancer. Implement Auto Scaling group for each ECS service by using the Cluster Autoscaler.

C. Configure the ECS services to use the blue/green deployment type and an Application Load Balancer. Implement an Auto Scaling group for each ECS service by using the Cluster Autoscaler.

D. Configure the ECS services to use the blue/green deployment type and an Application Load Balancer. Implement Service Auto Scaling for each ECS service.

Suggested Answer: *A*

Community vote distribution

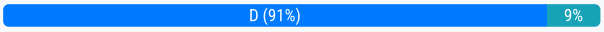

SmileyCloud Highly Voted 1 year ago

Selected Answer: D

A and B are out, it says the app uses HTTPS.
C is out because we have Fargate and there is no Cluster Auto Scaling there.
So, it's D because we have Service Auto Scaling. -> https://repost.aws/knowledge-center/ecs-fargate-service-auto-scaling

upvoted 15 times

emiliocb4 1 year ago

NLB supports HTTPS so why excluding A?

upvoted 1 times

SmileyCloud 12 months ago

Unlike a Classic Load Balancer or an Application Load Balancer, a Network Load Balancer can't have application layer (layer 7) HTTP or HTTPS listeners. It only supports transport layer (layer 4) TCP listeners. HTTP and HTTPS traffic can be routed to your environment over TCP.

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environments-cfg-nlb.html#

upvoted 9 times

ForDummies 1 month, 1 week ago

But this ELB will be used to on-demand video application, so ALB is out. Also, ECS will use HTTPS, not ELB.

upvoted 1 times

career360guru Most Recent 7 months, 1 week ago

Selected Answer: D

Option D

upvoted 1 times

ggrodskiy 11 months, 1 week ago

Correct D.

upvoted 1 times

Hypercuber 11 months, 1 week ago

Selected Answer: D

Answer is D. For those voting C, remember that it's on Fargate, so there is no such cluster autoscaling.

upvoted 4 times

nicecurls 11 months, 3 weeks ago

Selected Answer: D

select D. for Fargate there is no Cluster Auto Scaling there.

upvoted 2 times

NikkyDicky 11 months, 4 weeks ago

Selected Answer: D

D

no NLB for ECS, no Cluster for Fargate

upvoted 2 times

vjp_training 10 months, 2 weeks ago

D is correct but you can use NLB for ECS. Key word is Service Auto Scaling

https://docs.aws.amazon.com/AmazonECS/latest/userguide/create-network-load-balancer.html

upvoted 1 times

bhanus 1 year ago

Selected Answer: D

@MODERATOR, PLEASE remove my previous comment as I mentioned C.

As per comment from SmileyCloud , C is not correct because there is no Cluster Auto Scaling. D is the answer.

Thank you @SmileyCloud for clarifying

D is the answer

upvoted 2 times

SkyZeroZx 1 year ago

Selected Answer: D

https://repost.aws/knowledge-center/ecs-fargate-service-auto-scaling

upvoted 2 times

easytoo 1 year ago

d-d-d-d-d-d-d

upvoted 1 times

james55 1 year ago

Selected Answer: D

"Amazon ECS cluster auto scaling is only supported with Auto Scaling group capacity providers. For Amazon ECS workloads that are hosted on AWS Fargate, see AWS Fargate capacity providers."

upvoted 2 times

bhanus 1 year ago

Selected Answer: C

AB are eliminated because of NLB

C has Auto Scaling Group with Cluster Autoscaler: As per ChatGPT - By implementing an Auto Scaling group for each ECS service using the Cluster Autoscaler, you can automatically adjust the number of tasks (containers) based on the demand. The Cluster Autoscaler scales the ECS tasks in response to CloudWatch alarms, allowing you to scale the infrastructure up or down to handle the increasing number of users.

upvoted 3 times

bhanus 1 year ago

changing my vote to D as SmileyCloud pointed. for Fargate there is no Cluster Auto Scaling there.

upvoted 2 times

Question #224 *Topic 1*

A company is running a containerized application in the AWS Cloud. The application is running by using Amazon Elastic Container Service (Amazon ECS) on a set of Amazon EC2 instances. The EC2 instances run in an Auto Scaling group.

The company uses Amazon Elastic Container Registry (Amazon ECR) to store its container images. When a new image version is uploaded, the new image version receives a unique tag.

The company needs a solution that inspects new image versions for common vulnerabilities and exposures. The solution must automatically delete new image tags that have Critical or High severity findings. The solution also must notify the development team when such a deletion occurs.

Which solution meets these requirements?

A. Configure scan on push on the repository. Use Amazon EventBridge to invoke an AWS Step Functions state machine when a scan is complete for images that have Critical or High severity findings. Use the Step Functions state machine to delete the image tag for those images and to notify the development team through Amazon Simple Notification Service (Amazon SNS).

B. Configure scan on push on the repository. Configure scan results to be pushed to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Lambda function when a new message is added to the SQS queue. Use the Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Email Service (Amazon SES).

C. Schedule an AWS Lambda function to start a manual image scan every hour. Configure Amazon EventBridge to invoke another Lambda function when a scan is complete. Use the second Lambda function to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Notification Service (Amazon SNS).

D. Configure periodic image scan on the repository. Configure scan results to be added to an Amazon Simple Queue Service (Amazon SQS) queue. Invoke an AWS Step Functions state machine when a new message is added to the SQS queue. Use the Step Functions state machine to delete the image tag for images that have Critical or High severity findings. Notify the development team by using Amazon Simple Email Service (Amazon SES).

Suggested Answer: *C*

Community vote distribution

A (100%)

joleneinthebackyard Highly Voted 1 year, 2 months ago

Selected Answer: A

You want to look for "scan on push" solution, as scanning periodically is not enough, damage might have been done -> C, D is out, only A, B
A sounds complex, but B even worse, how can you put result in SQS? wording is so bad if they means sending message to SQS. Notifying by SES is a straight red flag that AWS exams like to use.
Only A makes sense.

upvoted 10 times

kz407 9 months, 1 week ago

Problem with this approach is, if you scan only what's pushed, and it has a zero-day vulnerability, you won't see it. Since you are scanning only when you are pushing, you won't detect the vulnerability ever. IMO, scanning periodically gives a better shot. Ideally it should be scanning both on push and periodically.

upvoted 3 times

kz407 Most Recent 9 months, 1 week ago

Selected Answer: A

https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning.html

In a nutshell, 2 types of scans.

Basic: Scanned against CVE DB, "ON PUSH" or a manual scan. Don't see any way of notifying anywhere.
Enhanced: Ongoing scanning with Amazon Inspector, findings delivered via EventBridge notifications.

Closest answer would be A.

upvoted 2 times

shaaam80 1 year, 1 month ago

Selected Answer: A

Answer A.

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: A

Option A

upvoted 2 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

A, but I think step function need to call Lambda to delete tag. there is not direct ecr integration

upvoted 3 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: A

Use the building feature if you can, so scan on push.

I go with A because other options are not good B - you cannot use SES.

upvoted 2 times

Maria2023 1 year, 6 months ago

Selected Answer: A

I vote A since I tested it and confirm it's achievable. As for B - I couldn't find any option to publish the result of the scan to SQS so I stopped there

upvoted 1 times

elanelans 1 year, 6 months ago

Selected Answer: A

A meet the requirements.

https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning.html

https://docs.aws.amazon.com/AmazonECR/latest/userguide/ecr-eventbridge.html

upvoted 2 times

SmileyCloud 1 year, 6 months ago

Selected Answer: A

C and D are out because they are not automatic but rather scheduled.

B is out because you don't need SQS for this and def don't need SES.

A makes sense because it's much leaner solution.

upvoted 2 times

nexus2020 1 year, 6 months ago

Selected Answer: A

Use the building feature if you can, so scan on push. And A make more sense

upvoted 1 times

bhanus 1 year, 6 months ago

Selected Answer: A

I go with A because other options are not good

B - you cannot use SES. SES is generally used to send Bulk/marketing emails.

C- schedule Lambda to scan every hour is not a good approach

D - like B you cannot use SES for this use case.

So A sounds reasonable

upvoted 2 times

emiliocb4 1 year, 6 months ago

why not A ?

upvoted 1 times

Question #225 *Topic 1*

A company runs many workloads on AWS and uses AWS Organizations to manage its accounts. The workloads are hosted on Amazon EC2. AWS Fargate. and AWS Lambda. Some of the workloads have unpredictable demand. Accounts record high usage in some months and low usage in other months.

The company wants to optimize its compute costs over the next 3 years. A solutions architect obtains a 6-month average for each of the accounts across the organization to calculate usage.

Which solution will provide the MOST cost savings for all the organization's compute usage?

A. Purchase Reserved Instances for the organization to match the size and number of the most common EC2 instances from the member accounts.

B. Purchase a Compute Savings Plan for the organization from the management account by using the recommendation at the management account level.

C. Purchase Reserved Instances for each member account that had high EC2 usage according to the data from the last 6 months.

D. Purchase an EC2 Instance Savings Plan for each member account from the management account based on EC2 usage data from the last 6 months.

Suggested Answer: *B*

Community vote distribution

B (100%)

elanelans Highly Voted 1 year, 6 months ago

Selected Answer: B

A. Incorrect: RI's Supports only EC2 instances.
B. Correct: Compute savings plan supports EC2, Fargate and Lambda. Applied in Organization's management account.
C. Incorrect: RI's Supports only EC2 instances and Changes to be applied at Organizations management account.
D. Incorrect: Instance Saving plan supports only EC2.

upvoted 14 times

titi_r Most Recent 8 months, 4 weeks ago

Selected Answer: B

B - "Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66%. These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, Region, OS or tenancy, and also apply to Fargate or Lambda usage.

https://aws.amazon.com/savingsplans/compute-pricing/

upvoted 1 times

shaaam80 1 year, 1 month ago

Answer B. Compute Savings plan covers EC2, Fargate & Lambda. Instance Savings plan only for EC2 instances.

upvoted 3 times

career360guru 1 year, 1 month ago

Selected Answer: B

Option B

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: B

its a B

upvoted 1 times

SkyZeroZx 1 year, 6 months ago

Selected Answer: B

A. Incorrect: RI's Supports only EC2 instances.
B. Correct: Compute savings plan supports EC2, Fargate and Lambda. Applied in Organization's management account.

C. Incorrect: RI's Supports only EC2 instances and Changes to be applied at Organizations management account.

D. Incorrect: Instance Saving plan supports only EC2.

upvoted 2 times

SmileyCloud 1 year, 6 months ago

Selected Answer: B

B, magic keywords - Management account and Compute savings Plan.

upvoted 1 times

nexus2020 1 year, 6 months ago

Selected Answer: B

Compute Savings plan is made for this usage type

upvoted 1 times

bhanus 1 year, 6 months ago

Selected Answer: B

B- compute savings plans covers all ec2, fargate, lambda.

upvoted 1 times

Question #226 *Topic 1*

A company has hundreds of AWS accounts. The company uses an organization in AWS Organizations to manage all the accounts. The company has turned on all features.

A finance team has allocated a daily budget for AWS costs. The finance team must receive an email notification if the organization's AWS costs exceed 80% of the allocated budget. A solutions architect needs to implement a solution to track the costs and deliver the notifications.

Which solution will meet these requirements?

A. In the organization's management account, use AWS Budgets to create a budget that has a daily period. Add an alert threshold and set the value to 80%. Use Amazon Simple Notification Service (Amazon SNS) to notify the finance team.

B. In the organization's management account, set up the organizational view feature for AWS Trusted Advisor. Create an organizational view report for cost optimization. Set an alert threshold of 80%. Configure notification preferences. Add the email addresses of the finance team.

C. Register the organization with AWS Control Tower. Activate the optional cost control (guardrail). Set a control (guardrail) parameter of 80%. Configure control (guardrail) notification preferences. Use Amazon Simple Notification Service (Amazon SNS) to notify the finance team.

D. Configure the member accounts to save a daily AWS Cost and Usage Report to an Amazon S3 bucket in the organization's management account. Use Amazon EventBridge to schedule a daily Amazon Athena query to calculate the organization's costs. Configure Athena to send an Amazon CloudWatch alert if the total costs are more than 80% of the allocated budget. Use Amazon Simple Notification Service (Amazon SNS) to notify the finance team.

Suggested Answer: *A*

Community vote distribution

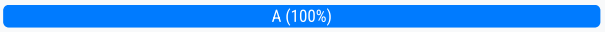

elanelans Highly Voted 2 years ago

Selected Answer: A

A. Makes sense.
B. Trusted advisor not required.
C. Control Tower not required.
D. Budgets can be managed in Org's Mgmt account itself.

upvoted 10 times

85b5b55 Most Recent 4 months, 1 week ago

Selected Answer: A

AWS Budgets can help for daily tracking and notify through SNS.

upvoted 1 times

kgpoj 10 months, 3 weeks ago

Selected Answer: A

A: AWS Budgets + SNS = Easy budget (daily) tracking and alerts

B: Trusted Advisor is for recommendations, not daily budgets.

C: Control Tower is for governance, not budget alerts

D: Complex setup with no added value over AWS Budgets

upvoted 2 times

career360guru 1 year, 7 months ago

Selected Answer: A

Option A

upvoted 1 times

nicecurls 1 year, 11 months ago

Selected Answer: A

ofc it's Ahttps://www.examtopics.com/exams/amazon/aws-certified-solutions-architect-professional-sap-c02/view/#

upvoted 3 times

NikkyDicky 1 year, 11 months ago

Selected Answer: A

straight A

upvoted 2 times

SkyZeroZx 1 year, 12 months ago

Selected Answer: A

A. Makes sense.

B. Trusted advisor not required.

C. Control Tower not required.

D. Budgets can be managed in Org's Mgmt account itself.

upvoted 2 times

rxhan 1 year, 11 months ago

you copy and paste other people answers

upvoted 6 times

easytoo 2 years ago

a-a-a-a-a-a-a

upvoted 1 times

SmileyCloud 2 years ago

Selected Answer: A

This one is simple. A

upvoted 1 times

nexus2020 2 years ago

Selected Answer: A

A, simple one

upvoted 1 times

MoussaNoussa 2 years ago

A is the answer

upvoted 1 times

bhanus 2 years ago

Selected Answer: A

A is the answer

upvoted 1 times

Question #227 *Topic 1*

A company provides auction services for artwork and has users across North America and Europe. The company hosts its application in Amazon EC2 instances in the us-east-1 Region. Artists upload photos of their work as large-size. high-resolution image files from their mobile phones to a centralized Amazon S3 bucket created in the us-east-1 Region. The users in Europe are reporting slow performance for their image uploads.

How can a solutions architect improve the performance of the image upload process?

A. Redeploy the application to use S3 multipart uploads.

B. Create an Amazon CloudFront distribution and point to the application as a custom origin.

C. Configure the buckets to use S3 Transfer Acceleration.

D. Create an Auto Scaling group for the EC2 instances and create a scaling policy.

Suggested Answer: *C*

Community vote distribution

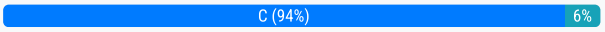

chico2023 Highly Voted 1 year, 4 months ago

Selected Answer: C

Main point of the question: "The users in Europe are reporting slow performance for their image uploads."
How do we improve performance? If we look on the latency side, sure, S3 Transfer Acceleration (option C), but the question puts another variable to our scenario: "Artists upload photos of their work as large-size. high-resolution image files from their mobile phones..."
If you just look at that above, you would switch to A as we can improve upload with multipart.

Here comes the plot twist "The users in Europe are reporting slow performance for their image uploads." - Meaning, in "Europe", not in the "NA". Of course! The bucket in the US... So yeah, question really bad, not objective (in my pov) and with lots of interpretations, but C would help them with the perception of performance in this context.

upvoted 29 times

Jay_2pt0_1 1 year ago

Kudos to you for such a great explanation!

upvoted 2 times

kpcert Most Recent 6 months, 3 weeks ago

Selected Answer: C

Between A and C, I would choose C - Transfer Acceleration, as this issue is focusing on improving the upload performance across the region

upvoted 1 times

rohan0411 6 months, 3 weeks ago

Why not B ?

upvoted 1 times

Monsterpuss 2 weeks, 4 days ago

Because Cloudwatch is a CDN aimed ad delivering out content, not uploading it.

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: C

Option C. As the users in Europe only are facing this issue. A would improve upload performance overall for both US and Europe.

upvoted 3 times

Pupu86 1 year, 1 month ago

I believe this question should rightfully be a multi-choice question where A and C are the answer together to solve this problem statement

https://aws.amazon.com/blogs/compute/uploading-large-objects-to-amazon-s3-using-multipart-upload-and-transfer-acceleration/

upvoted 2 times

skyhiker 1 year, 4 months ago

I would choose A. Why does C say "Configure the buckets [more than one] to use S3 Transfer Acceleration?
Sometimes you have to hate how these questions and answers are worded.
upvoted 1 times

skyhiker 1 year, 4 months ago
C would be the answer if the 's' was removed. Will to go with C.
upvoted 1 times

RGR21 1 year, 5 months ago
Selected Answer: A
I have some doubts about this question, it makes more sense to use multipart upload to split the file and gain upload speed. AWS Transfer Accelerator seems to be applied to reduce delay.http
s://aws.amazon.com/pt/blogs/compute/uploading-large-objects-to-amazon-s3-using-multipart-upload-and-transfer-acceleration/
upvoted 1 times

ggrodskiy 1 year, 5 months ago
Correct C.
upvoted 1 times

NikkyDicky 1 year, 5 months ago
Selected Answer: C
C
would be good in combination with A, but better as a standalone choice
upvoted 1 times

Christina666 1 year, 5 months ago
Selected Answer: C
upload performance-> transfer acceleration
upvoted 1 times

javitech83 1 year, 6 months ago
Selected Answer: C
correct is C
upvoted 1 times

pupsik 1 year, 6 months ago
Selected Answer: A
Transfer Acceleration doesn't guarantee a significant increase in upload speed.
A multi-part upload on other hand does, because it uploads multiple smaller chunks of the files in parallel.

Ideally multi-part upload and Transfer Accelerator should be deployed together. If we had to pick only one of the two, multi-part upload would result in better performance.

https://aws.amazon.com/blogs/compute/uploading-large-objects-to-amazon-s3-using-multipart-upload-and-transfer-acceleration/
upvoted 1 times

YodaMaster 1 year, 5 months ago
Using your link, the tests mentioned show C is faster
Single upload with transfer acceleration 40% faster
Multipart upload without transfer acceleration 38% faster
upvoted 4 times

SeemaDataReader 11 months, 1 week ago
Reading carefully into the blog looks like the author did some maths wrong.
Multipart upload took 43s which is 40% faster than base of 72s
Transfer acceleration took 45s which is 38% faster than base of 72s.
So based on this multipart gives better performance
upvoted 1 times

shmoeee 4 months, 2 weeks ago
Double check your math brother..
upvoted 1 times

SkyZeroZx 1 year, 6 months ago
Selected Answer: C

C. https://aws.amazon.com/s3/transfer-acceleration/

upvoted 1 times

SmileyCloud 1 year, 6 months ago

Selected Answer: C

C. https://aws.amazon.com/s3/transfer-acceleration/

upvoted 2 times

MoussaNoussa 1 year, 6 months ago

C of course

upvoted 1 times

bhanus 1 year, 6 months ago

Selected Answer: C

C - Transfer acceleration. S3 Transfer Acceleration utilizes the Amazon CloudFront global network of edge locations to accelerate the transfer of data to and from S3 buckets. By enabling S3 Transfer Acceleration on the centralized S3 bucket, the users in Europe will experience faster uploads as their data will be routed through the closest CloudFront edge location.

upvoted 1 times

Question #228 *Topic 1*

A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web. application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.

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

A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tiers. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SQS).

B. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session data. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.

C. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node groups. Use ReplicaSets to run the web servers and applications. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system across all EKS pods to store frontend web server session data.

D. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node groups. Run the web servers and application as Kubernetes deployments in the EKS cluster. Store the frontend web server session data in an Amazon DynamoDB table. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.

Suggested Answer: *B*

Community vote distribution

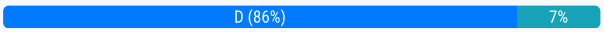

pupsik Highly Voted 2 years ago

Selected Answer: D

A looked good until "store session data in SQS".

upvoted 24 times

SkyZeroZx Highly Voted 2 years ago

Selected Answer: D

what a worst ques

A - Why do you need SQS to store web sever session data. SQS is for decoupling services

B - EBS multi attach is for SAME availibility zone. The ques says multipel availibility zones

C - Why do you need EFS to store web sever session data. Its damn expensive

D - Better answer- But again why need for EKS.

If I were to choose one option, its D as its better compared to ABC

upvoted 16 times

JoeTromundo Most Recent 8 months, 3 weeks ago

Selected Answer: D

By exclusion of the other options, the least worst answer is D.

upvoted 3 times

liuliangzhou 9 months, 3 weeks ago

Selected Answer: D

B,D can do it all.

Multiple EC2 accesses a single storage using EFS and EBS, with priority given to EFS file storage.

DynamoDB can also be used for session storage.

https://docs.aws.amazon.com/aws-sdk-php/v2/guide/feature-dynamodb-session-handler.html

upvoted 1 times

43c89f4 1 year, 2 months ago

one of the poor Question... so answer we give poor... its D. because i cant choose ABC

upvoted 4 times

ayadmawla 1 year, 6 months ago

Selected Answer: B

I think that the issue of multi-attach EBS in one AZ is dealt with by the manner in which it is explained. It is the EC2 that are distributed in Multi-AZ not the EBS. Just my pov.

upvoted 3 times

career360guru 1 year, 7 months ago

Selected Answer: D

Option D, Though C is also possible but Multi-attach EBS has higher operational overhead.

upvoted 2 times

covabix879 1 year, 9 months ago

Selected Answer: D

Due to operational efficiency D is better choice compared to B.

upvoted 1 times

task_7 1 year, 9 months ago

Selected Answer: D

deployments carry ReplicaSets

DynamoDB table for session data

upvoted 1 times

rsn 1 year, 9 months ago

Selected Answer: C

There is a requirement for fault tolerance. I feel 'C" satisfies that as it has replicasets. Option D does not talk about it

upvoted 1 times

skyhiker 1 year, 10 months ago

Now i'll have to go with B. Check out what alabiba says to question, "Can aws sqs be used to store web server session data?"

alabiba "No, AWS SQS (Simple Queue Service) is not typically used for storing web server session data. SQS is a message queuing service that is designed for reliable and scalable message communication between distributed systems. For storing session data, it is more common to use dedicated session storage solutions such as databases (e.g., Amazon DynamoDB) or in-memory caches (e.g., Redis)."

upvoted 2 times

chikorita 1 year, 9 months ago

problem with option B is " Multi-Attach on EC2 instances that are distributed across multiple Availability Zones"; please note that multi-attach can only span since AZ

option D is correct

upvoted 2 times

NikkyDicky 1 year, 11 months ago

Selected Answer: D

D - best of the worst

upvoted 7 times

YodaMaster 1 year, 11 months ago

Selected Answer: D

A looked good until "store session data in SQS".

upvoted 2 times

Henrytml 1 year, 12 months ago

A looked good until "store session data in SQS".

upvoted 3 times

javitech83 2 years ago

Selected Answer: D

A looked good until "store session data in SQS".

upvoted 2 times

Maria2023 2 years ago

Selected Answer: A

Fargate is the service, the only question remains the storage. Amazon EBS Multi-Attach is single-az service, so remains A. Even though I am not very confident with SQS caching web service sessions.

upvoted 1 times

PhuocT 2 years ago

Agree, this is a worst question

D is best choice for this question, but I would prefer to change EKS to ECS Fargate for compute and Elasticache for Redis for session.

upvoted 3 times

Question #229 *Topic 1*

A solutions architect is planning to migrate critical Microsoft SQL Server databases to AWS. Because the databases are legacy systems, the solutions architect will move the databases to a modern data architecture. The solutions architect must migrate the databases with near-zero downtime.

Which solution will meet these requirements?

A. Use AWS Application Migration Service and the AWS Schema Conversion Tool (AWS SCT). Perform an in-place upgrade before the migration. Export the migrated data to Amazon Aurora Serverless after cutover. Repoint the applications to Amazon Aurora.

B. Use AWS Database Migration Service (AWS DMS) to rehost the database. Set Amazon S3 as a target. Set up change data capture (CDC) replication. When the source and destination are fully synchronized, load the data from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB instance.

C. Use native database high availability tools. Connect the source system to an Amazon RDS for Microsoft SQL Server DB instance. Configure replication accordingly. When data replication is finished, transition the workload to an Amazon RDS for Microsoft SQL Server DB instance.

D. Use AWS Application Migration Service. Rehost the database server on Amazon EC2. When data replication is finished, detach the database and move the database to an Amazon RDS for Microsoft SQL Server DB instance. Reattach the database and then cut over all networking.

Suggested Answer: *C*

Community vote distribution

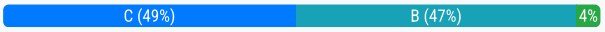

SmileyCloud Highly Voted 2 years ago

Selected Answer: C

C. The proper way is to use AWS DMS, but the answer here uses S3 (???) which will take forever. So the answer is C.

upvoted 17 times

yorkicurke 1 year, 7 months ago

the following link maybe helpful for some;
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.S3.html#CHAP_Target.S3.Limitations

upvoted 3 times

Ganshank Highly Voted 1 year, 10 months ago

C
https://aws.amazon.com/blogs/database/part-3-migrating-to-amazon-rds-for-sql-server-using-transactional-replication-with-native-backup-and-restore/

upvoted 12 times

Kaps443 Most Recent 2 weeks, 5 days ago

Selected Answer: C

Uses native replication; minimal downtime; direct migration

upvoted 1 times

0b43291 7 months, 1 week ago

Selected Answer: C

By using native database high availability tools and replication methods, you can achieve near-zero downtime during the migration process. The other options may not provide the same level of seamless data replication and minimal downtime as the native SQL Server replication tools.

Option B: Use AWS Database Migration Service (AWS DMS) to rehost the database. Set Amazon S3 as a target. Set up change data capture (CDC) replication. When the source and destination are fully synchronized, load the data from Amazon S3 into an Amazon RDS for Microsoft SQL Server DB instance.

While AWS DMS can be used for migrations, it introduces additional complexity compared to native SQL Server replication tools.
Staging data in Amazon S3 and then loading into the target RDS instance can cause downtime during the final cutover.
Native replication tools can directly replicate data to the target RDS instance without an intermediate storage solution.

upvoted 4 times

jefnmet 10 months, 2 weeks ago

B look correct to me
upvoted 1 times

helloworldabc 10 months, 1 week ago
just C
upvoted 1 times

8693a49 11 months ago
Selected Answer: B
B is the AWS way of doing a DB migration. C might work and could be better in some cases, but because the DBs are legacy you might run into compatibility issues or limitations with the tooling. If the databases are very large you really want to use B because you need to ship the bulk of the data with Snowball.
upvoted 1 times

CAIYasia 11 months, 1 week ago
Selected Answer: C
C. Correct, Near-Zero Downtime.
A. In-Place Upgrade and Migration to Aurora: This involves multiple steps and the potential for increased downtime during the cutover process.
Schema Conversion: Depending on the complexity of the legacy system, converting schemas and ensuring compatibility with Amazon Aurora can be challenging and time-consuming.
B. Intermediate Storage in Amazon S3: Adds complexity
Two-Step Process: First replicating to Amazon S3 and then loading into Amazon RDS adds additional steps and potential points of failure.
upvoted 1 times

grandcanyon 11 months, 4 weeks ago
Selected Answer: B
In option C - "Connect the source system to an Amazon RDS for Microsoft SQL Server DB instance", should be target, not source
upvoted 2 times

trungtd 1 year ago
Selected Answer: C
Use AWS Database Migration Service (AWS DMS) to "rehost" the database????
How you can "rehost" database with DMS
upvoted 1 times

michele_scar 1 year, 1 month ago
Selected Answer: B
DMS should be the better service for this use case
upvoted 2 times

titi_r 1 year, 1 month ago
Selected Answer: C
Answer: C.
upvoted 1 times

BrijMohan08 1 year, 2 months ago
Selected Answer: A
AWS Application Migration Service (MGN) is a highly automated lift-and-shift (rehost) solution that simplifies the migration of applications to AWS. It supports near-zero downtime migrations by continuously replicating the source servers to AWS.

Repointing the applications to Amazon Aurora Serverless satisfies the migration to the modern data architecture.
upvoted 2 times

svenkata18 1 year, 3 months ago
Why not A as the question the it should rearchitected from legacy
upvoted 1 times

JOKERO 1 year, 3 months ago
Native database high availability (HA) tools include the Always On or distributed availability group clusters in Microsoft SQL Server and Oracle's Data Guard replications. This approach requires a major effort to set up across extended, cross-site HA clusters, and might cause some performance degradation because of the longer latency to achieve fully synchronous active/active deployments. However, this method provides the closest to near-zero downtime during the cutover.
upvoted 4 times

ftaws 1 year, 5 months ago

What is "native database high availability tools"????

upvoted 1 times

tmlong18 1 year, 5 months ago

Selected Answer: C

B. Use AWS Database Migration Service (AWS DMS) to rehost the database.

This action is not 'rehost'

upvoted 1 times

adelynlllllllll 1 year, 6 months ago

C:

Use distributed AG, it will work.

https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-sql-server-to-aws-using-distributed-availability-groups.html

upvoted 1 times

Question #230 *Topic 1*

A company's solutions architect is analyzing costs of a multi-application environment. The environment is deployed across multiple Availability Zones in a single AWS Region. After a recent acquisition, the company manages two organizations in AWS Organizations. The company has created multiple service provider applications as AWS PrivateLink-powered VPC endpoint services in one organization. The company has created multiple service consumer applications in the other organization.

Data transfer charges are much higher than the company expected, and the solutions architect needs to reduce the costs. The solutions architect must recommend guidelines for developers to follow when they deploy services. These guidelines must minimize data transfer charges for the whole environment.

Which guidelines meet these requirements? (Choose two.)

A. Use AWS Resource Access Manager to share the subnets that host the service provider applications with other accounts in the organization.

B. Place the service provider applications and the service consumer applications in AWS accounts in the same organization.

C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.

D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name.

E. Create a Savings Plan that provides adequate coverage for the organization's planned inter-Availability Zone data transfer usage.

Suggested Answer: *AB*

Community vote distribution

SkyZeroZx Highly Voted 1 year, 12 months ago

Selected Answer: AD

A By sharing the subnets that host the service provider applications using AWS Resource Access Manager (RAM), the service consumer applications can be deployed in the same organization's accounts. This allows the traffic between the service consumer and service provider applications to stay within the organization's network, reducing data transfer charges.
D By using the Availability Zone-specific endpoint service's local DNS name, the service consumer compute resources can directly access the service provider applications within the same Availability Zone. This eliminates the need for cross-Availability Zone data transfer, thus reducing data transfer charges.

upvoted 14 times

helloworldabc 10 months, 1 week ago

just CD

upvoted 1 times

xav1er Highly Voted 1 year, 9 months ago

Selected Answer: CD

- **C. Turn off cross-zone load balancing for the Network Load Balancer in all service provider application deployments.**
- **D. Ensure that service consumer compute resources use the Availability Zone-specific endpoint service by using the endpoint's local DNS name.**

upvoted 8 times

youonebe Most Recent 7 months ago

Selected Answer: CD

Normally when data leaves AZ to another, there is a cost associated

upvoted 1 times

sam2ng 7 months, 4 weeks ago

This is why C is correct:

"For ALB and CLB, there is no cross-AZ data transfer charges within the same VPC. But for NLB, if the client and target are in one AZ, but the NLB is in another AZ, there will be a zone-in and zone-out which is $0.02."

upvoted 2 times

JoeTromundo 8 months, 3 weeks ago

Selected Answer: CD

B is not an option: While placing resources in the same organization might simplify management, it does not inherently reduce data transfer charges. Data transfer costs between AWS Organizations accounts are typically not impacted by being in the SAME OR DIFFERENT organizations, especially when using PrivateLink.

upvoted 1 times

vip2 11 months, 3 weeks ago

Selected Answer: CD

C D is correct one

For C, Cross-zone load balancing can distribute traffic across multiple AZs, which increases data transfer costs between AZs. Disabling cross-zone load balancing ensures that traffic remains within the same AZ, reducing the associated data transfer charges. This is particularly important for applications using AWS PrivateLink, as it will help keep data transfers within the same AZ as much as possible.

upvoted 1 times

michele_scar 1 year, 1 month ago

Selected Answer: CD

B is useless because if you place the resource in the same org but in different AZs you will pay the same as different org in different AZs. So B is uncorrect (like A and E).

Remains C and D as a solution that should reduce costs.

upvoted 3 times

seetpt 1 year, 1 month ago

Selected Answer: BD

BD for me

upvoted 2 times

4555894 1 year, 2 months ago

B - allows data transfer between linked accounts to be free of charge.

D - ensures traffic stays within the same AZ as much as possible, minimizing inter-AZ data transfer costs.

CD - Save money.

upvoted 1 times

VerRi 1 year, 3 months ago

Selected Answer: BD

"The company manages two organisations in AWS Organizations," which means they have one organisation for service providers and one more for consumers.

A. Since applications are created in the provider organisation, sharing the subnet with other accounts within the same organisation has no effect.

B. Combining provider and consumer into one organisation is the first move for Option D.

C. Cross-zone load balancing does not change the amount of data traffic passing through the NLB, it affects how that traffic is distributed across the targets.

D. AZ-specific endpoint helps to reduce data transfer charges because it keeps the traffic in a single AZ and is designed for intra-regional communication within the same account or organization.

E. WTF

upvoted 4 times

Dgix 1 year, 3 months ago

Selected Answer: BD

It's B and D.

A. Sharing subnets does not directly reduce data transfer charges.

C. Turning off cross-zone load balancing does not impact data transfer costs between VPC endpoints and service consumers.

E. A Savings Plan reduces costs for compute usage, not specifically for data transfer charges.

upvoted 5 times

mav3r1ck 1 year, 3 months ago

Turning off cross-zone load balancing can reduce inter-AZ data transfer costs. With cross-zone load balancing disabled, a Network Load Balancer (NLB) only routes requests to targets in the same Availability Zone as the load balancer node that received the request. This setup reduces the data transferred across Availability Zones, thereby reducing costs.

upvoted 3 times

ajeeshb 1 year, 3 months ago

Selected Answer: CD

Answer: C, D

upvoted 3 times

marszalekm 1 year, 4 months ago

https://docs.aws.amazon.com/ram/latest/userguide/shareable.html "Can share with only AWS accounts in its own organization." ec2:Subnet

upvoted 2 times

Wardove 1 year, 4 months ago

Selected Answer: CD

Answer is CD

D) Obvious option, This approach minimizes data transfer costs by ensuring that traffic between service consumers and service providers stays within the same Availability Zone

C) Only after setting up your NLB, you can create a VPC Endpoint Service (VPC-E) that is powered by AWS PrivateLink. Cross-zone lb feature is optional for NLB since 2018 so, turning off cross-zone load balancing can help ensure that data does not unnecessarily cross Availability Zones, thereby once again reducing data transfer costs

https://docs.aws.amazon.com/elasticloadbalancing/latest/userguide/how-elastic-load-balancing-works.html

B) Incorrect: putting the workloads into 1 org - would not make any effect on billing neither, unless you change the topology profoundly and move away the VPCE solution - but we are not talking about Re-architecting, we are looking to provide guidelines

A) Incorrect: RAM can be used only within 1 organization

E) Incorrect: there is no a such flavor of Saving plans, AWS provides 3 Compute, EC Instance and SageMaker Saving plans

upvoted 6 times

JOKERO 1 year, 3 months ago

You can also share with specific AWS accounts by account ID, regardless of whether the account is part of an organization.

upvoted 1 times

LazyAutonomy 1 year, 5 months ago

Selected Answer: BD

Holy bageezus, never seen a discussion thread so divided.

@NikkyDicky is spot on - cross zone traffic is indeed where the money is going. I think we all know that.

A - appears incorrect, we cannot share subnets between accounts in different AWS Orgs. Even if you could, or even if you chose A+B, it would be impractical to assume all other workloads could be deployed in service provider subnets. Would probably run out of IPs. And even if the subnets were huge and we didn't run out of IPs, there is no mechanism in A to guide developers deploying their workloads to reduce or prevent cross-AZ traffic. You could share the subnets and deploy all provider/consumer workloads in the same set of subnets and still end up with the same huge bill :-)

upvoted 5 times

LazyAutonomy 1 year, 5 months ago

B - appears correct. @Just_Ninja's explanation nails it. If you use Organizations and you create accounts, then in each member account, the logical identifiers for each availability zone (e.g. "eu-central-1a") are guaranteed to map to the same AZ Physical ID (e.g. "euc1-az3") for all accounts within the Organization. In other words, it's likely that AZ "eu-central-1a" for accounts in OrgABC is not the same as AZ "eu-central-1a" for accounts in OrgXYZ. That's a problem if you're trying to eliminate unnecessary cross-zone traffic. Without this, you could instruct developers to use AZ-specific DNS names and still end up with the same huge bill :-)

upvoted 1 times

LazyAutonomy 1 year, 5 months ago

C - appears incorrect, but the reason has nothing to do with "compromising high availability". As pointed out by @elmoh, cross-zone load balancing isn't enabled by default in NLBs anyway. See https://docs.aws.amazon.com/elasticloadbalancing/latest/network/network-load-balancers.html#cross-zone-load-balancing. Even if cross-zone load balancing was enabled by default in NLBs, this option doesn't cover the Gateway Load Balancer VPC endpoint service use case.

upvoted 2 times

tmlong18 1 year, 5 months ago

Selected Answer: CD

I go with C & D.

Data transfer cost base on physical distance.(cross AZ, cross region, internal)

A & B - shared VPC doesn't distribute traffic to inter-az

upvoted 3 times

Jay_2pt0_1 1 year, 6 months ago

This question is poorly framed. I go with A & D, not because they are great, but because the others are terrible. You should not have to move into the same org (that can't be the answer). Also, we won't compromise HA, so that can't be the answer either.

upvoted 3 times

Question #231 *Topic 1*

A company has an on-premises Microsoft SQL Server database that writes a nightly 200 GB export to a local drive. The company wants to move the backups to more robust cloud storage on Amazon S3. The company has set up a 10 Gbps AWS Direct Connect connection between the on-premises data center and AWS.

Which solution meets these requirements MOST cost-effectively?

A. Create a new S3 bucket. Deploy an AWS Storage Gateway file gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share.

B. Create an Amazon FSx for Windows File Server Single-AZ file system within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system. Enable nightly backups.

C. Create an Amazon FSx for Windows File Server Multi-AZ file system within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to an SMB file share on the Amazon FSx file system. Enable nightly backups.

D. Create a new S3 bucket. Deploy an AWS Storage Gateway volume gateway within the VPC that is connected to the Direct Connect connection. Create a new SMB file share. Write nightly database exports to the new SMB file share on the volume gateway, and automate copies of this data to an S3 bucket.

Suggested Answer: *A*

Community vote distribution

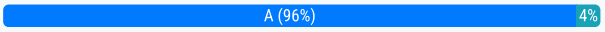

SkyZeroZx Highly Voted 1 year, 5 months ago

Selected Answer: A

File Gateway == SMB , NFS

Volumes Gateway == iSCSI

Tape Gateway = VTL

upvoted 32 times

SIJUTHOMASP Most Recent 6 months ago

Selected Answer: A

Guys, options B and C are exactly same. :)

upvoted 1 times

duriselvan 1 year ago

Ans D

he most cost-effective solution for moving the backups to S3 is D. Deploy an AWS Storage Gateway volume gateway, create an SMB file share, and automate data copies to S3.

Here's why:

Cost-effectiveness: Volume gateways use Amazon EBS volumes for local storage, which is typically more cost-effective than Amazon FSx for Windows File Server for storing large amounts of data. Additionally, this approach avoids the need for additional backups within Amazon FSx, further reducing costs.

Direct Connect utilization: Leveraging the existing Direct Connect connection optimizes network bandwidth for transferring data to S3, minimizing latency and potential data transfer charges.

Automated backups: Automating copies of the nightly exports to S3 ensures reliable backups and minimizes manual intervention.

upvoted 1 times

career360guru 1 year, 1 month ago

Selected Answer: A

Option A

upvoted 2 times

yorkicurke 1 year, 1 month ago

Selected Answer: A

if you read the end of the following link's paragraph, its right there in documentation;
https://aws.amazon.com/storagegateway/features/#Gateway_Types
under " Amazon S3 File Gateway "
Customers can use Amazon S3 File Gateway to back up on-premises file data as objects in Amazon S3 (including Microsoft SQL Server and Oracle databases and logs), and for hybrid cloud workflows using data generated by on-premises applications for processing by AWS services such as machine learning or big data analytics.

upvoted 2 times

NolaHOla 1 year, 1 month ago

But the answer A never mentions S3 file gateway?

upvoted 2 times

SK_Tyagi 1 year, 4 months ago

Selected Answer: A

https://aws.amazon.com/storagegateway/features/

upvoted 1 times

rafael796 1 year, 4 months ago

Selected Answer: A

file gateway = most cheap

upvoted 1 times

NikkyDicky 1 year, 5 months ago

Selected Answer: A

A - SMB mount = file gwy

upvoted 2 times

RockyLeon 1 year, 6 months ago

Selected Answer: A

file gateway -> used to store file inside s3
volume gateway -> used to store file in on-premises using iSCSI connectivity

upvoted 2 times

Jackhemo 1 year, 6 months ago

Using Olabiba.ai to learn not to find an answer:

Jack: Labiba, what is the Microsoft SQL Server database export is it block or file?

oLabiba: The Microsoft SQL Server database export is typically a block-level backup. It captures the data at the database level, including the schema, tables, and records, and stores it in a binary format. This allows for efficient backup and restoration of the database.

In summary, if you primarily need file-level access to your backups, File Gateway is a better choice. If you require block-level storage and want to optimize for low-latency access, Volume Gateway is a better fit.

Let me know if you know the answer now.

upvoted 2 times

Maria2023 1 year, 6 months ago

Selected Answer: A

File Gateway could be mapped as SMB file share and used by the database or other automation to transfer database backups. Volume Gateway is more used to perform volume snapshots on the on-premise system so I don't believe it's a sustainable approach here.

upvoted 3 times

SmileyCloud 1 year, 6 months ago

Selected Answer: A

It's A (file gateway). Volume gateway is iSCSI.

upvoted 2 times

Jackhemo 1 year, 6 months ago

Selected Answer: D

olabiba.ai says D
Option D: Using an AWS Storage Gateway volume gateway allows you to write the nightly database exports to an SMB file share on the volume

gateway, which can be stored locally and automatically backed up to an S3 bucket. This solution is cost-effective as it utilizes the existing Direct Connect connection and requires minimal additional infrastructure.
upvoted 2 times

easytoo 1 year, 6 months ago
d-d-d-d-d-d

By deploying an AWS Storage Gateway volume gateway within the VPC connected to the Direct Connect connection, the company can leverage the high-speed, low-latency connection to transfer the nightly database exports to the SMB file share on the volume gateway. This allows for efficient and reliable data transfer.

Automating copies of this data from the SMB file share to an S3 bucket provides a cost-effective solution for storing the backups in more robust cloud storage on Amazon S3. The company can take advantage of the durability, scalability, and cost-effectiveness of S3 for long-term storage.
upvoted 2 times

nexus2020 1 year, 6 months ago
Selected Answer: A
Between A and D:

write to local drive can also be a network drive mapped to the windows server. therefore SME file share is enough (A), D is Block level, for sure will cost more.

the File Gateway is designed for file-level access and presents Amazon S3 storage as a file share, while the Volume Gateway provides block-level access and appears as local block storage volumes. The choice between the two depends on the specific needs and requirements of your applications and data access patterns.
upvoted 2 times

bjexamprep 1 year ago
The backend of storage gateway is actually S3 storage, which means both volume gateway and file gateway share the same cost for storage. And the gateway cost is the same according to aws pricing: https://aws.amazon.com/storagegateway/pricing/.
so where did you get the "D is Block level, for sure will cost more"?
upvoted 1 times

bhanus 1 year, 6 months ago
I am between A and D. ChatGpt says A. But The reason why I think D is because, the question says backups are written to local drive(which means its a volume on onpremises machine). So I thought a volume can be attached to volume gateway. But ChatGPT says In terms of cost-effectiveness and simplicity, option A is a better choice. It involves using an AWS Storage Gateway file gateway, which directly stores the data as objects in Amazon S3 without the need for on-premises storage. This eliminates the complexity and costs associated with maintaining an on-premises volume gateway.
upvoted 1 times

bhanus 1 year, 6 months ago
I might be wrong with my theory. Going with A
upvoted 1 times

Jackhemo 1 year, 6 months ago
Use olabiba.ai. It is better.
upvoted 1 times

PhuocT 1 year, 6 months ago
Q: are you using openAI as your AI engine?
olabiba.com: Yes, I am powered by OpenAI's advanced AI technology. It allows me to understand and respond to your messages in a conversational manner. OpenAI provides the foundation for my capabilities, but the Olabiba team has also customized and trained me to better suit your needs. So, feel free to ask me anything or share your thoughts!
upvoted 2 times

gd1 1 year, 6 months ago
Volume will iSCSI so hat is out. Therefor A is correct
upvoted 1 times

Question #232 *Topic 1*

A company needs to establish a connection from its on-premises data center to AWS. The company needs to connect all of its VPCs that are located in different AWS Regions with transitive routing capabilities between VPC networks. The company also must reduce network outbound traffic costs, increase bandwidth throughput, and provide a consistent network experience for end users.

Which solution will meet these requirements?

A. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPC. Create VPC peering connections that initiate from the central VPC to all other VPCs.

B. Create an AWS Direct Connect connection between the on-premises data center and AWS. Provision a transit VIF, and connect it to a Direct Connect gateway. Connect the Direct Connect gateway to all the other VPCs by using a transit gateway in each Region.

C. Create an AWS Site-to-Site VPN connection between the on-premises data center and a new central VPUse a transit gateway with dynamic routing. Connect the transit gateway to all other VPCs.

D. Create an AWS Direct Connect connection between the on-premises data center and AWS. Establish an AWS Site-to-Site VPN connection between all VPCs in each Region. Create VPC peering connections that initiate from the central VPC to all other VPCs.

Suggested Answer: *B*

Community vote distribution

Pupu86 Highly Voted 1 year, 7 months ago

Selected Answer: B

In fact site to site VPN would be more affordable than deploying a Direct Connect leased line. However, AWS also wants to market their product by stating that there is a need to increase throughput (site to site only can achieve max of 1.25Gbps) and consistent user experience (AWS Direct Connect > Site-to-Site VPN) so B would be a better choice.

upvoted 9 times

gfhbox0083 Most Recent 11 months, 3 weeks ago

B, for sure.

For a consistent network experience

upvoted 1 times

TonytheTiger 1 year, 3 months ago

Selected Answer: B

https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-gateway.html

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: B

Option B may not be most cost-effective best option in terms of performance.

upvoted 3 times

joleneinthebackyard 1 year, 8 months ago

Anyone can explain that why Site to Site VPN not valid?

upvoted 1 times

fartosh 1 year, 1 month ago

The company wants to increase bandwidth throughput, which is gained by establishing Direct Connect.

upvoted 2 times

Gabehcoud 1 year, 10 months ago

what if the situation is 1 AWS account, different VPC's across different regions? Can we still use a TGW?

upvoted 1 times

hexie 1 year, 11 months ago

Selected Answer: B

B.

Cant be D because TGW doesnt support transitive connections, so if users connect to a VPN it invalidate this options.

A and C are skippable on the first phrase.

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: B

B no doubt

upvoted 1 times

SkyZeroZx 2 years ago

Selected Answer: B

direct connect + vpc = direct connect gw + TGW. so B

upvoted 3 times

rxhan 1 year, 11 months ago

Mr. copy and paste

upvoted 3 times

Maria2023 2 years ago

Selected Answer: B

Transit gateway is a regional service but you can peer different TGs in different regions

https://aws.amazon.com/about-aws/whats-new/2019/12/aws-transit-gateway-supports-inter-region-peering/

upvoted 1 times

SmileyCloud 2 years ago

Selected Answer: B

B. No need for D and S2S VPN.

upvoted 1 times

aragon_saa 2 years ago

BBBBBBBBBB?

upvoted 1 times

nexus2020 2 years ago

Selected Answer: B

direct connect + vpc = direct connect gw + TGW. so B

upvoted 3 times

Question #233 *Topic 1*

A company is migrating its development and production workloads to a new organization in AWS Organizations. The company has created a separate member account for development and a separate member account for production. Consolidated billing is linked to the management account. In the management account, a solutions architect needs to create an IAM user that can stop or terminate resources in both member accounts.

Which solution will meet this requirement?

A. Create an IAM user and a cross-account role in the management account. Configure the cross-account role with least privilege access to the member accounts.

B. Create an IAM user in each member account. In the management account, create a cross-account role that has least privilege access. Grant the IAM users access to the cross-account role by using a trust policy.

C. Create an IAM user in the management account. In the member accounts, create an IAM group that has least privilege access. Add the IAM user from the management account to each IAM group in the member accounts.

D. Create an IAM user in the management account. In the member accounts, create cross-account roles that have least privilege access. Grant the IAM user access to the roles by using a trust policy.

Suggested Answer: *D*

Community vote distribution

D (100%)

bhanus Highly Voted 2 years ago

Selected Answer: D

D - Cross account role should be created in destination(member) account. The role has trust entity to master account.

upvoted 6 times

duriselvan Most Recent 1 year, 6 months ago

A is ans

A. Create an IAM user and a cross-account role in the management account. Configure the cross-account role with least privilege access to the member accounts.

Here's why:

Cross-account roles: Provide a secure and managed way for users or services in one AWS account to access resources in another account.
Least privilege access: Configure the cross-account role with the minimum permissions needed to stop or terminate resources in the member accounts, minimizing potential security risks.
Centralized control: Maintaining user credentials and access in the management account simplifies centralized management and auditing.

upvoted 1 times

helloworldabc 10 months, 1 week ago

just D

upvoted 1 times

career360guru 1 year, 7 months ago

Selected Answer: D

Option D

upvoted 2 times

skyhiker 1 year, 10 months ago

Hmm, seems like alot of work. What if the question was, In the management account, a solutions architect needs to create an IAM user that can stop or terminate resources in 100 organization or member accounts? Asked AI "Using AWS Organizations, can you create both IAM user and permission sets in the management account for accessing managed organization resources?" The answer was Yes.

upvoted 1 times

NikkyDicky 1 year, 11 months ago

Selected Answer: D

its a D