

높은 통과율 AWS-Solutions-Associate 최고 품질 덤프 공부 자료 인증 시험 공부

EMC DEP-3CR1 PowerProtect Cyber Recovery Exam 3

질문 # 26
 An enterprise customer needs a Cyber Recovery solution to be implemented. As an outcome from a previous workshop, the following backup environment needs to be protected to the CR Vault.
 Location 1: 4 PowerProtect DDs
 Location 2: 4 PowerProtect DDs
 Location 3: 2 PowerProtect DDs
 Location 4: 2 PowerProtect DDs
 The customer wants to implement a CR Vault in a 5th location.
 How many Cyber Recovery systems must be installed at a minimum level?

- A. 0
- B. 1
- C. 2
- D. 3

정답 B

질문 # 27
 What vault status is displayed if none of the PowerProtect DD systems in the CR Vault are able to communicate with the Cyber Recovery software?

- A. Unknown
- B. Unlocked
- C. Locked
- D. Degraded

정답 B

질문 # 28

DEP-3CR1 높은 통과율 시험대비 덤프 공부:
https://www.koreadumps.com/DEP-3CR1_exam-braindumps.html

- DEP-3CR1 시험문제집 DEP-3CR1 시험대비 덤프 다운로드 DEP-3CR1 시험덤프 무료 다운로드를 위해 지금 www.itdumps.com 에서 DEP-3CR1 검색 DEP-3CR1 높은 통과율 인기 덤프 문제
- DEP-3CR1 최신버전 덤프 공부 DEP-3CR1 높은 통과율 인기 덤프 문제 DEP-3CR1 시험대비 덤프 다운로드 지금 www.itdumps.com 에서 DEP-3CR1 를 검색하고 무료로 다운로드하세요 DEP-3CR1 높은 통과율 인기 덤프 문제
- DEP-3CR1 덤프 문제 DEP-3CR1 시험패스 가능 덤프 DEP-3CR1 시험대비 공부하기 검색만 하면 www.itdumps.com 에서 DEP-3CR1 * 무료 다운로드 DEP-3CR1 최신기출자료
- 최신버전 DEP-3CR1 시험대비 덤프 최신자료 덤프 문제 검색만 하면 www.itdumps.com 에서 DEP-3CR1 < 무료 다운로드 DEP-3CR1 시험패스 가능한 인증 공부 자료
- DEP-3CR1 최신버전 덤프 공부 DEP-3CR1 합격보장 가능 덤프 공부 DEP-3CR1 시험대비 덤프 다운로드 지금 www.itdumps.com (유니) 인기 무료 다운로드를 위해 DEP-3CR1 를 검색하십시오 DEP-3CR1 시험패스 가능 덤프
- DEP-3CR1 최신버전 덤프 공부 DEP-3CR1 시험대비 공부하기 DEP-3CR1 덤프 문제

DEP-3CR1 시험대비 덤프 최신자료 & DEP-3CR1 높은 통과율 시험대비 덤프 공부

참고: KoreaDumps에서 Google Drive로 공유하는 무료, 최신 AWS-Solutions-Associate 시험 문제집이 있습니다:
<https://drive.google.com/open?id=1qIc-6BlpWfZXbQxOuU4s0xAH6mVgOLxa>

KoreaDumps는 Amazon 인증 AWS-Solutions-Associate 시험에 대하여 가이드를 해줄 수 있는 사이트입니다. KoreaDumps는 여러분의 전문지식을 업그레이드 시켜줄 수 있고 또한 한번에 Amazon 인증 AWS-Solutions-Associate 시험을 패스하도록 도와주는 사이트입니다. KoreaDumps 제공하는 자료들은 모두 업계 전문가들이 자신의 지식과 끈임없는 경험 등으로 만들어낸 퍼펙트 자료들입니다. 품질은 정확도 모두 보장되는 문제집입니다. Amazon 인증 AWS-Solutions-Associate 시험은 여러분이 이 지식을 한층 업할 수 있는 시험이며 우리 또한 일년 무료 업데이트 서비스를 제공합니다.

우리의 덤프는 기존의 시험문제와 답과 시험문제 분석 등입니다. KoreaDumps에서 제공하는 Amazon AWS-Solutions-Associate 시험 자료의 문제와 답은 실제 시험의 문제와 답과 아주 비슷합니다. KoreaDumps는 여러분이 한번에 Amazon AWS-Solutions-Associate 인증 시험을 패스함을 보장 드립니다.

>> AWS-Solutions-Associate 최고 품질 덤프 공부 자료 <<

AWS-Solutions-Associate 최고 품질 덤프 공부 자료 시험은 저희 최신 덤프로 패스가능

KoreaDumps는 여러분이 원하는 최신 최고 버전의 Amazon 인증 AWS-Solutions-Associate 덤프를 제공합니다. Amazon

인증AWS-Solutions-Associate덤프는 IT업계전문가들이 끊임없는 노력과 지금까지의 경험으로 연구하여 만들어낸 제일 정확한 시험문제와 답들로 만들어졌습니다. KoreaDumps의 문제집으로 여러분은 충분히 안전이 시험을 패스하실 수 있습니다. 우리 KoreaDumps 의 문제집들은 모두 100%합격율을 자랑하며 KoreaDumps의 제품을 구매하였다면 Amazon 인증AWS-Solutions-Associate시험패스와 자격증 취득은 근심하지 않으셔도 됩니다. 여러분은 IT업계에 서 또 한층 업그레이드 될것입니다.

최신 AWS Certified Solutions Architect AWS-Solutions-Associate 무료 샘플 문제 (Q567-Q572):

질문 # 567

A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve individual newspaper pages via a website written in Java They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate Its archive to AW5 and produce a cost efficient architecture and still be designed for availability and durability. Which is the most appropriate?

- A. Model the environment using CloudFormation use an EC2 instance running Apache webserver and an open source search application, stripe multiple standard EB5 volumes together to store the JPEGs and search index.
- B. Use a CloudFront download distribution to serve the JPEGs to the end users and Install the current commercial search product, along with a Java Container Tor the website on EC2 instances and use Route53 with DNS round-robin.
- C. Use a single-AZ RD5 MySQL instance lo store the search index 33d the JPEG images use an EC2 instance to serve the website and translate user queries into 5QL.
- **D. Use S3 with standard redundancy to store and serve the scanned files, use CloudSearch for query processing, and use Elastic Beanstalk to host the website across multiple availability zones.**
- E. Use S3 with reduced redundancy lo store and serve the scanned files, install the commercial search application on EC2 Instances and configure with auto-scaling and an Elastic Load Balancer.

정답: D

설명:

There is no such thing as "Most appropriate" without knowing all your goals. I find your scenarios very fuzzy, since you can obviously mix-n-match between them. I think you should decide by layers instead:

Load Balancer Layer: ELB or just DNS, or roll-your-own. (Using DNS+EIPs is slightly cheaper, but less reliable than ELB.)

Storage Layer for 17TB of Images: This is the perfect use case for S3. Off-load all the web requests directly to the relevant JPEGs in S3. Your EC2 boxes just generate links to them.

If your app already serves it's own images (not links to images), you might start with EFS. But more than likely, you can just setup a web server to re-write or re-direct all JPEG links to S3 pretty easily.

If you use S3, don't serve directly from the bucket- Serve via a CNAME in domain you control. That way, you can switch in CloudFront easily.

EBS will be way more expensive, and you'll need 2x the drives if you need 2 boxes. Yuck. Consider a smaller storage format. For example, JPEG200 or WebP or other tools might make for smaller images.

There is also the DejaVu format from a while back.

Cache Layer: Adding Cloud Front in front of S3 will help people on the other side of the world-- well, possibly. Typical archives follow a power law. The long tail of requests means that most JPEGs won't be requested enough to be in the cache. So you are only speeding up the most popular objects. You can always wait, and switch in CF later after you know your costs better. (In some cases, it can actually lower costs.)

You can also put CloudFront in front of your app, since your archive search results should be fairly static.

This will also allow you to run with a smaller instance type, since CF will handle much of the load if you do it right.

Database Layer: A few options:

Use whatever your current server does for now, and replace with something else down the road.

Don't under-estimate this approach, sometimes it's better to start now and optimize later.

Use RDS to run MySQL/ Postgres

I'm not as familiar with ElasticSearch I Cloudsearch, but obviously Cloudsearch will be less maintenance+setup.

App Layer:

When creating the app layer from scratch, consider Cloud Formation and/or OpsWorks. It's extra stuff to learn, but helps down the road.

Java+ Tomcat is right up the alley ofElasticBeanstalk. (Basically EC2 + Autoscale + ELB).

Preventing Abuse: When you put something in a public S3 bucket, people will hot-link it from their web pages. If you want to prevent that, your app on the EC2 box can generate signed links to S3 that expire in a few hours. Now everyone will be forced to go thru the app, and the app can apply rate limiting, etc.

Saving money: If you don't mind having downtime:

run everything in one AZ (both DBs and EC2s). You can always add servers and AZs down the road, as long as it's architected to be stateless. In fact, you should use multiple regions if you want it to be really robust.

use Reduced Redundancy in S3 to save a few hundred bucks per month (Someone will have to "go fix it" every time it breaks, including having an off-line copy to repair S3.)

Buy Reserved Instances on your EC2 boxes to make them cheaper. (Start with the RI market and buy a partially used one to get started.) It's just a coupon saying "if you run this type of box in this AZ, you will save on the per-hour costs." You can get 1/2 to 1/3 off easily.

Rewrite the application to use less memory and CPU -that way you can run on fewer/ smaller boxes.

(May or may not be worth the investment.)

If your app will be used very infrequently, you will save a lot of money by using Lambda. I'd be worried that it would be quite slow if you tried to run a Java application on it though ..

We're missing some information like load, latency expectations from search, indexing speed, size of the search index, etc. But with what you've given us, I would go with S3 as the storage for the files (S3 rocks. It is really, really awesome). If you're stuck with the commercial search application, then on EC2 instances with autoscaling and an ELB. If you are allowed an alternative search engine, Elasticsearch is probably your best bet. I'd run it on EC2 instead of the AWS Elasticsearch service, as IMHO it's not ready yet.

Don't autoscale Elasticsearch automatically though, it'll cause all sorts of issues. I have zero experience with CloudSearch so I can't comment on that. Regardless of which option, I'd use Cloud Formation for all of it.

질문 # 568

A company has data collection sensors at different locations. The data collection sensors stream a high volume of data to the company. The company wants to design a platform on AWS to ingest and process high-volume streaming data. The solution must be scalable and support data collection in near real time. The company must store the data in Amazon S3 for future reporting. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Lambda to deliver streaming data and store the data to Amazon S3.
- B. Use AWS Database Migration Service (AWS DMS) to deliver streaming data to Amazon S3.
- C. Use AWS Glue to deliver streaming data to Amazon S3.
- **D. Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3.**

정답: D

설명:

To ingest and process high-volume streaming data with the least operational overhead, Amazon Kinesis Data Firehose is a suitable solution. Amazon Kinesis Data Firehose can capture, transform, and deliver streaming data to Amazon S3 or other destinations. Amazon Kinesis Data Firehose can scale automatically to match the throughput of the data and handle any amount of data. Amazon Kinesis Data Firehose is also a fully managed service that does not require any servers to provision or manage.

References:

What Is Amazon Kinesis Data Firehose?

Amazon Kinesis Data Firehose Pricing

질문 # 569

A Solutions Architect is architecting a workload that requires a performant object-based storage system that must be shared with multiple Amazon EC2 instances.

Which AWS service meets this requirement?

- A. Amazon EBS
- **B. Amazon S3**
- C. Amazon EFS
- D. Amazon ElastiCache

정답: B

질문 # 570

A gaming company is developing a game that requires significant compute resources to process game logic, player interactions, and real-time updates. The company needs a compute solution that can dynamically scale based on fluctuating player demand while maintaining high performance. The company must use a relational database that can run complex queries.

- A. Deploy an Amazon Elastic Container Service (Amazon ECS) cluster on AWS Fargate to supply compute capacity. Scale the ECS tasks based on player demand. Use Amazon Aurora Serverless v2 as the database.
- B. Refactor the game logic into small, stateless functions. Use AWS Lambda to process the game logic. Use Amazon DynamoDB as the database.
- C. Use AWS ParallelCluster for high performance computing (HPC). Provision compute nodes that have GPU instances to process the game logic and player interactions. Use Amazon RDS for MySQL as the database.
- D. Deploy Amazon EC2 instances to supply compute capacity. Configure Auto Scaling groups to achieve dynamic scaling based on player count. Use Amazon RDS for MySQL as the database.

정답: A

설명:

Amazon ECS with AWS Fargate provides serverless container compute that can scale tasks dynamically in response to demand. This matches the unpredictable scaling requirements of a gaming workload. Aurora Serverless v2 provides on-demand, autoscaling relational database capacity while supporting complex SQL queries. EC2 with Auto Scaling (A) works but requires significant management overhead. Lambda with DynamoDB (B) is not suitable because the workload requires a relational database and long-running processes. ParallelCluster with HPC (D) is designed for batch scientific workloads, not dynamic, interactive gaming. Option C provides the correct balance of elasticity, performance, and managed services for both compute and relational database needs. References:* Amazon ECS User Guide - Scaling ECS tasks with Fargate* Amazon Aurora User Guide - Aurora Serverless v2 for dynamic workloads

질문 # 571

A company stores its application logs in an Amazon CloudWatch Logs log group. A new policy requires the company to store all application logs in Amazon OpenSearch Service (Amazon Elasticsearch Service) in near-real time. Which solution will meet this requirement with the LEAST operational overhead?

- A. Install and configure Amazon Kinesis Agent on each application server to deliver the logs to Amazon Kinesis Data Streams. Configure Kinesis Data Streams to deliver the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service)
- B. Create an Amazon Kinesis Data Firehose delivery stream. Configure the log group as the delivery stream's source. Configure Amazon OpenSearch Service (Amazon Elasticsearch Service) as the delivery stream's destination.
- C. Create an AWS Lambda function. Use the log group to invoke the function to write the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).
- D. Configure a CloudWatch Logs subscription to stream the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).

정답: C

설명:

Explanation

<https://computingforgeeks.com/stream-logs-in-aws-from-cloudwatch-to-elasticsearch/>

질문 # 572

.....

KoreaDumps 에서는 Amazon AWS-Solutions-Associate 시험에 대비한 고품질 Amazon AWS-Solutions-Associate덤프를 제공해 드립니다. Amazon AWS-Solutions-Associate덤프는 IT 업계 종사자들에 있어서 아주 중요한 인증시험이자 인기 자격증을 취득할 수 있는 필수과목입니다. Amazon AWS-Solutions-Associate시험을 합격하여 자격증을 취득하시면 취업하는데 가산점이 될 수 있고 급여 인상이나 이직에도 많은 도움을 드릴 수 있습니다.

AWS-Solutions-Associate최신 업데이트버전 인증덤프 : https://www.koreadumps.com/AWS-Solutions-Associate_exam-braindumps.html

KoreaDumps 표 Amazon인증AWS-Solutions-Associate덤프를 공부하시면 시험보는데 자신감이 생기고 시험불합격에 대한 우려도 줄어들 것입니다, 업데이트서비스 제공, IT인증자격증을 취득할 수 있는 AWS-Solutions-Associate인기시험을 어떻게 패스할가 고민그만하시고 저희 IT전문가들이 제작한 AWS-Solutions-Associate 인기덤프자료를 데려가주세요, Amazon AWS-Solutions-Associate최고품질 덤프공부자료 Software 버전은 테스트용으로 PDF 버전 공부를 마친후 시험전에 실력테스트 가능합니다, Amazon인증 AWS-Solutions-Associate시험을 패스하여 자격증을 취득하면 소원이 이루어집니다, 그 중 Amazon AWS-Solutions-Associate인증시험을 패스한 분들도 모두 KoreaDumps인증시험덤프를 사용하였기 때문입니다.

id=1qle-6BlpWfZXbQxOuU4s0xAH6mVgOLxa