# Amazon.SAP-C02.v2025-06-02.q210

| | |
|---|---|
| **Exam Code:** | SAP-C02 |
| **Exam Name:** | AWS Certified Solutions Architect - Professional (SAP-C02) |
| **Certification Provider:** | Amazon |
| **Free Question Number:** | 210 |
| **Version:** | v2025-06-02 |
| **# of views:** | 107 |
| **# of Questions views:** | 2100 |
| https://www.freepdfdumps.com/Amazon.SAP-C02.v2025-06-02.q210.html | |

**NEW QUESTION: 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.
**Answer: C (LEAVE A REPLY)**
https://aws.amazon.com/migration-evaluator/features/

**NEW QUESTION: 2**
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 Single Sign-On (AWS SSO) to implement this functionality.
Which solution will meet these requirements MOST cost-effectively?
**A.** Create an organization in AWS Organizations. Turn on the AWS SSO 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 AWS SSO 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 AWS SSO feature in Organizations Create and configure an AD Connector to connect to the company's on-premises Active Directory. Configure AWS SSO 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 AWS SSO 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 ADdirectory.
**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 AWS SSO 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.
**Answer: D (LEAVE A REPLY)**
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_org_support-all-features.html
https://docs.aws.amazon.com/singlesignon/latest/userguide/get-started-prereqs-considerations.html

**NEW QUESTION: 3**
A company has purchased appliances from different vendors. The appliances all have loT 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 loT sensors to AWS loT 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. UseAmazon Athena and Amazon OuickSight for analysis.
**B.** Migrate the application server to AWS Fargate, which will receive the information from loT sensors and parse the information into a relational format. Save the parsed information to Amazon Redshift 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.
**Answer: A (LEAVE A REPLY)**
* 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.
This solution involves connecting the IoT sensors to the AWS IoT Core, setting a rule to invoke an AWS Lambda function to parse the information, and saving a .csv file to Amazon S3. AWS Glue can be used to catalog the files and Amazon Athena and Amazon QuickSight can be used for analysis. This solution will enable faster and more cost-effective data analysis.
This solution is in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that: "AWS IoT Core can be used to ingest and process the data, AWS Lambda can be used to process and transform the data, and Amazon S3 can be used to store the data. AWS Glue can be used to catalog and access the data, Amazon Athena can be used to query the data, and Amazon QuickSight can be used to visualize the data." (Source:https://d1.awsstatic.com
/training-and-certification/docs-sa-pro
/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf)

**NEW QUESTION: 4**
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.
**Answer: C (LEAVE A REPLY)**
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.

**NEW QUESTION: 5**
An ecommerce company runs an application on AWS. The application has an Amazon API Gateway API that invokes an AWS Lambda function. The data is stored in an Amazon RDS for PostgreSQL DB instance.
During the company's most recent flash sale, a sudden increase in API calls negatively affected the application's performance. A solutions architect reviewed the Amazon CloudWatch metrics during that time and noticed a significant increase in Lambda invocations and database connections. The CPU utilization also was high on the DB instance.
What should the solutions architect recommend to optimize the application's performance?
**A.** Increase the memory of the Lambda function. Modify the Lambda function to close the database connections when the data is retrieved.
**B.** Add an Amazon ElastiCache for Redis cluster to store the frequently accessed data from the RDS database.
**C.** Create an RDS proxy by using the Lambda console. Modify the Lambda function to use the proxy endpoint.
**D.** Modify the Lambda function to connect to the database outside of the function's handler. Check for an existing database connection before creating a new connection.
**Answer: C (LEAVE A REPLY)**
This option will optimize the application's performance by reducing the overhead of opening and closing database connections for each Lambda invocation. An RDS proxy is a fully managed database proxy for Amazon RDS that makes applications more scalable, more resilient to database failures, andmore secure1. It allows applications to pool and share connections established with the database, improving database efficiency and application scalability1. By creating an RDS proxy by using the Lambda console, you caneasily configure your Lambda function to use the proxy endpoint instead of the directdatabase endpoint2.

This will enable your Lambda function to reuse existing connections from the proxy's connection pool, reducing the latency and CPU utilization caused byestablishing new connections for each invocation. It will also prevent connection saturation or exhaustion on the database, which can degrade performance or cause errors3.

**NEW QUESTION: 6**
A car rental company has built a serverless REST API to provide data to its mobile app. The app consists of an Amazon API Gateway API with a Regional endpoint, AWS Lambda functions, and an Amazon Aurora MySQL Serverless DB cluster. The company recently opened the API to mobile apps of partners. A significant increase in the number of requests resulted, causing sporadic database memory errors. Analysis of the API traffic indicates that clients are making multiple HTTP GET requests for the same queries in a short period of time. Traffic is concentrated during business hours, with spikes around holidays and other events.
The company needs to improve its ability to support the additional usage while minimizing the increase in costs associated with the solution.
Which strategy meets these requirements?
**A.** Convert the API Gateway Regional endpoint to an edge-optimized endpoint. Enable caching in the production stage.
**B.** Implement an Amazon ElastiCache for Redis cache to store the results of the database calls. Modify the Lambda functions to use the cache.
**C.** Modify the Aurora Serverless DB cluster configuration to increase the maximum amount of available memory.
**D.** Enable throttling in the API Gateway production stage. Set the rate and burst values to limit the incoming calls.
**Answer: (SHOW ANSWER)**
Explanation: This option allows the company to use Amazon CloudFront to improve the latency and availability of the API requests by caching the responses at the edge locations closest to the clients1. By enabling caching in the production stage, the company can reduce the number of calls made to the backend services, such as Lambda functions and Aurora Serverless DB cluster, and save on costs and resources2. This option also helps to handle traffic spikes and reduce database memory errors by serving cached responses instead of querying the database repeatedly.
:
Choosing an API endpoint type
Enabling API caching to enhance responsiveness

**NEW QUESTION: 7**
A company has deployed its database on an Amazon RDS for MySQL DB instance in the us-east-1 Region.
The company needs to make its data available to customers in Europe. The customers in Europe must have access to the same data as customers in the United States (US) and will not tolerate high

application latency or stale data. The customers in Europe and the customers in the US need to write to the database. Both groups of customers need to see updates from the other group in real time.
Which solution will meet these requirements?
**A.** Create an Amazon Aurora MySQL replica of the RDS for MySQL DB instance. Pause application writes to the RDS DB instance. Promote the Aurora Replica to a standalone DB cluster. Reconfigure the application to use the Aurora database and resume writes. Add eu-west-1 as a secondary Region to the 06 cluster. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu- west-1.
**B.** Add a cross-Region replica in eu-west-1 for the RDS for MySQL DB instance. Configure the replica to replicate write queries back to the primary DB instance. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
**C.** Copy the most recent snapshot from the RDS for MySQL DB instance to eu-west-1. Create a new RDS for MySQL DB instance in eu-west-1 from the snapshot. Configure MySQL logical replication from us- east-1 to eu-west-1. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the RDS for MySQL endpoint in eu-west-1.
**D.** Convert the RDS for MySQL DB instance to an Amazon Aurora MySQL DB cluster. Add eu-west-1 as a secondary Region to the DB cluster. Enable write forwarding on the DB cluster. Deploy the application in eu-west-1. Configure the application to use the Aurora MySQL endpoint in eu-west-1.
**Answer: D (LEAVE A REPLY)**
The company should use AWS Amplify to create a static website for uploads of media files. The company should use Amplify Hosting to serve the website through Amazon CloudFront. The company should use Amazon S3 to store the uploaded media files. The company should use Amazon Cognito to authenticate users.
This solution will meet the requirements with the least operational overhead because AWS Amplify is a complete solution that lets frontend web and mobile developers easily build, ship, and host full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as use cases evolve. No cloud expertise needed1. By using AWS Amplify, the company can refactor the application to a serverless architecture that reduces operational complexity and costs. AWS Amplify offers the following features and benefits:
* Amplify Studio: A visual interface that enables you to build and deploy a full-stack app quickly, including frontend UI and backend.
* Amplify CLI: A local toolchain that enables you to configure and manage an app backend with just a few commands.
* Amplify Libraries: Open-source client libraries that enable you to build cloud-powered mobile and web apps.
* Amplify UI Components: Open-source design system with cloud-connected components for building feature-rich apps fast.
* Amplify Hosting: Fully managed CI/CD and hosting for fast, secure, and reliable static and server-side rendered apps.
By using AWS Amplify to create a static website for uploads of media files, the company can leverage Amplify Studio to visually build a pixel-perfect UI and connect it to a cloud backend in clicks. By using

Amplify Hosting to serve the website through Amazon CloudFront, the company can easily deploy its web app or website to the fast, secure, and reliable AWS content delivery network (CDN), with hundreds of points of presence globally. By using Amazon S3 to store the uploaded media files, the company can benefit from a highly scalable, durable, and cost-effective object storage service that can handle any amount of data2. By using Amazon Cognito to authenticate users, the company can add user sign-up, sign-in, and access control to its web app with a fully managed service that scales to support millions of users3.

The other options are not correct because:

* Using AWS Application Migration Service to migrate the application server to Amazon EC2 instances would not refactor the application or accelerate development. AWS Application Migration Service (AWS MGN) is a service that enables you to migrate physical servers, virtual machines (VMs), or cloud servers from any source infrastructure to AWS without requiring agents or specialized tools.
However, this would not address the challenges of overutilization and data uploads failures. It would also not reduce operational overhead or costs compared to a serverless architecture.
* Creating a static website for uploads of media files and using AWS AppSync to create an API would not be as simple or fast as using AWS Amplify. AWS AppSync is a service that enables you to create flexible APIs for securely accessing, manipulating, and combining data from one or more data sources.
However, this would require more configuration and management than using Amplify Studio and Amplify Hosting. It would also not provide authentication features like Amazon Cognito.
* Setting up AWS IAM Identity Center (AWS Single Sign-On) to give users the ability to sign in to the application would not be as suitable as using Amazon Cognito. AWS Single Sign-On (AWS SSO) is a service that enables you to centrally manage SSO access and user permissions across multiple AWS accounts and business applications. However, this service is designed for enterprise customers who need to manage access for employees or partners across multiple resources. It is not intended for authenticating end users of web or mobile apps.

References:

https://aws.amazon.com/amplify/
https://aws.amazon.com/s3/
https://aws.amazon.com/cognito/
https://aws.amazon.com/mgn/
https://aws.amazon.com/appsync/
https://aws.amazon.com/single-sign-on/

**NEW QUESTION: 8**

A company wants to migrate its website from an on-premises data center onto AWS. At the same time, it wants to migrate the website to a containerized microservice-based architecture to improve the availability and cost efficiency. The company's security policy states that privileges and network permissions must be configured according to best practice, using least privilege.
A Solutions Architect must create a containerized architecture that meets the security requirements and has deployed the application to an Amazon ECS cluster.
What steps are required after the deployment to meet the requirements? (Choose two.)

**A.** Create tasks using the bridge network mode.
**B.** Create tasks using the awsvpc network mode.
**C.** Apply security groups to Amazon EC2 instances, and use IAM roles for EC2 instances to access other resources.
**D.** Apply security groups to the tasks, and pass IAM credentials into the container at launch time to access other resources.
**E.** Apply security groups to the tasks, and use IAM roles for tasks to access other resources.
**Answer: B,E (LEAVE A REPLY)**
Explanation: The awsvpc network mode provides each task with its own elastic network interface (ENI) and a primary private IP address1. By using this network mode, the solutions architect can isolate the tasks from each other and apply security groups to the tasks directly2. This way, the solutions architect can control the inbound and outbound traffic at the task level and enforce the least privilege principle3. IAM roles for tasks allow the solutions architect to assign permissions to each task separately, so that they can access other AWS resources that they need4. By using IAM roles for tasks, the solutions architect can avoid passing IAM credentials into the container at launch time, which is less secure and more prone to errors5.
:
awsvpc network mode
Task networking with the awsvpc network mode
Security groups for your VPC
IAM roles for tasks
Best practices for managing AWS access keys

**NEW QUESTION: 9**
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 (or Windows File Server file system. The tile 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-lime 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.
**Answer: C (LEAVE A REPLY)**
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html

**NEW QUESTION: 10**
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.
**Answer: (SHOW ANSWER)**
By creating an AMI that has Grafana pre-installed and storing the existing dashboards in Amazon Elastic File System (Amazon EFS) it allows for faster and more efficient scaling, and by creating an Auto Scaling group that uses the new AMI and setting the Auto Scaling group's minimum, desired, and maximum number of instances to one and creating an Application Load Balancer that serves at least two Availability Zones, it ensures high availability and minimized downtime.

**NEW QUESTION: 11**
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 withsecure, 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.
**Answer: B (LEAVE A REPLY)**
This solution will allow the external auditors to have read-only access to the company's AWS account while being compliant with AWS security best practices. By creating an IAM role, which is a secure and flexible way of granting access to AWS resources, and trusting the auditors' AWS account, the company can ensure that the auditors only have the permissions that are required for their role and nothing more. Assigning a unique external ID to the role's trust policy, it will ensure that only the auditors' AWS account can assume the role.
Reference:
AWS IAM Roles documentation:https://aws.amazon.com/iam/features/roles/
AWS IAM Best practices:https://aws.amazon.com/iam/security-best-practices/

**NEW QUESTION: 12**
A company wants to migrate its on-premises application to AWS. The database for the application stores structured product data and temporary user session data. The company needs to decouple the product data from the user session data. The company also needs to implement replication in another AWS Region for disaster recovery.
Which solution will meet these requirements with the HIGHEST performance?
**A.** Create an Amazon RDS DB instance to host the product data. Configure a read replica for the DB instance in another Region. Create an Amazon DynamoDB global table to host the user session data
**B.** Create an Amazon RDS DB instance to host the product data. Configure a read replica for the DB instance in another Region. Create a global datastore in Amazon ElastiCache for Memcached to host the user session data.
**C.** Create an Amazon RDS DB instance with separate schemas to host the product data and the user session data. Configure a read replica for the DB instance in another Region.
**D.** Create two Amazon DynamoDB global tables. Use one global table to host the product data Use the other global table to host the user session data. Use DynamoDB Accelerator (DAX) for caching.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 13**
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 lo 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 master 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 master account to create monthly reports for each business unit.
**C.** Configure AWS Budgets in each account and configure budget alerts lhat 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 master 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.
**Answer: B (LEAVE A REPLY)**
Configure AWS Budgets in the organization#€™s master 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 master account to create monthly reports for each business unit.
https://aws.amazon.com/about-aws/whats-new/2019/07/introducing-aws-budgets-reports/#:~:text=AWS%20Budgets%20gives%20you%20the,below%20the%20threshold%20you%20define

**NEW QUESTION: 14**
A company runs an unauthenticated static website (www.example.com) that includes a registration form for users. The website uses Amazon S3 for hosting and uses Amazon CloudFront as the content delivery network with AWS WAF configured. When the registration form is submitted, the website calls an Amazon API Gateway API endpoint that invokes an AWS Lambda function to process the payload and forward the payload to an external API call.
During testing, a solutions architect encounters a cross-origin resource sharing (CORS) error. The solutions architect confirms that the CloudFront distribution origin has the Access-Control-Allow-Origin header set towww.example.com.
What should the solutions architect do to resolve the error?
**A.** Change the CORS configuration on the S3 bucket. Add rules for CORS to the Allowed Origin element forwww.example.com.
**B.** Enable the CORS setting in AWS WAF. Create a web ACL rule in which the Access-Control-Allow-Origin header is set towww.example.com.

**C.** Enable the CORS setting on the API Gateway API endpoint. Ensure that the API endpoint is configured to return all responses that have the Access-Control -Allow-Origin header set towww.example.com.
**D.** Enable the CORS setting on the Lambda function. Ensure that the return code of the function has the Access-Control-Allow-Origin header set towww.example.com.
**Answer: C (LEAVE A REPLY)**
CORS errors occur when a web page hosted on one domain tries to make a request to a server hosted on another domain. In this scenario, the registration form hosted on the static website is trying to make a request to the API Gateway API endpoint hosted on a different domain, which is causing the error. To resolve this error, the Access-Control-Allow-Origin header needs to be set to the domain from which the request is being made. In this case, the header is already set to www.example.com on the CloudFront distribution origin.
Therefore, the solutions architect should enable the CORS setting on the API Gateway API endpoint and ensure that the API endpoint is configured to return all responses that have the Access-Control-Allow-Origin header set to www.example.com. This will allow the API endpoint to respond to requests from the static website without a CORS error.
https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-cors-errors/

**NEW QUESTION: 15**
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? (Select TWO.)
**A.** Deploy the security tool on EC2 instances in 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.
**Answer: A,D (LEAVE A REPLY)**
Option A, Deploy the security tool on EC2 instances in a new Auto Scaling group in the existing VPC, allows the company to use its existing security tool while still running it within the AWS environment. This ensures that all packets coming in and out of the VPC are inspected by the security tool in real time.
Option D, Provision a Gateway Load Balancer for each Availability Zone to redirect the traffic to the security tool, allows for high availability within an AWS Region. By provisioning a Gateway Load Balancer

for each Availability Zone, the traffic is redirected to the security tool in the event of any failures or outages. This ensures that the security tool is always available to inspect the traffic, even in the event of a failure.

**NEW QUESTION: 16**
A company runs a software-as-a-service <SaaS) application on AWS The application consists of AWS Lambda functions and an Amazon RDS for MySQL Multi-AZ database During market events the application has a much higher workload than normal Users notice slow response times during the peak periods because of many database connections The company needs to improve the scalable performance and availability of the database Which solution meets these requirements'?
**A.** Create an Amazon CloudWatch alarm action that triggers a Lambda function to add an Amazon RDS for MySQL read replica when resource utilization hits a threshold
**B.** Migrate the database to Amazon Aurora, and add a read replica Add a database connection pool outside of the Lambda handler function
**C.** Migrate the database to Amazon Aurora and add a read replica Use Amazon Route 53 weighted records
**D.** Migrate the database to Amazon Aurora and add an Aurora Replica Configure Amazon RDS Proxy to manage database connection pools
**Answer: D (LEAVE A REPLY)**
* Migrate to Amazon Aurora:
* Amazon Aurora is a MySQL-compatible, high-performance database designed to provide higher throughput than standard MySQL. Migrating the database to Aurora will enhance the performance and scalability of the database, especially under heavy workloads.
* Add Aurora Replica:
* Aurora Replicas provide read scalability and improve availability. Adding an Aurora Replica allows read operations to be distributed, thereby reducing the load on the primary instance and improving response times during peak periods.
* Configure Amazon RDS Proxy:
* Amazon RDS Proxy acts as an intermediary between the application and the Aurora database, managing connection pools efficiently. RDS Proxy reduces the overhead of opening and closing database connections, thus maintaining fewer active connections to the database and handling surges in database connections from the Lambda functions more effectively.
* This configuration reduces the database's resource usage and improves its ability to handle high volumes of concurrent connections.
References
* AWS Database Blog on RDS Proxy(Amazon Web Services, Inc.).
* AWS Compute Blog on Using RDS Proxy with Lambda(Amazon Web Services, Inc.).

**NEW QUESTION: 17**
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 adependency on a MySQL database A solutions architect must design a scalable and highly available solution to meet the demand of
200000 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
**Answer: (SHOW ANSWER)**
Using AWS CloudFormation to launch a stack with an Application Load Balancer (ALB) in front of an Amazon EC2 Auto Scaling group spanning three Availability Zones, a Multi-AZ deployment of an Amazon Aurora MySQL DB cluster with a Retain deletion policy, and an Amazon Route 53 alias record to route traffic from the company's domain to the ALB will ensure that

**NEW QUESTION: 18**
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 runninginstances 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.
**Answer: (SHOW ANSWER)**
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.

**NEW QUESTION: 19**
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.
**Answer: B (LEAVE A REPLY)**
The company should use AWS IoT Greengrass to deploy the ML model to the local server and provide local feedback to the factory workers. AWS IoT Greengrass is a service that extends AWS cloud capabilities to local devices, allowing them to collect and analyze data closer to the source of information,react autonomously to local events, and communicate securely with each other on local networks1. AWS IoT Greengrass also supports MLinference at the edge, enabling devices to run ML models locally without requiring internet connectivity2.
The other options are not correct because:
* Setting up an Amazon Kinesis video stream from each IP camera to AWS would not work if the factory's internet connectivity is down. It would also incur unnecessary costs and latency to stream video data to the cloud and back.
* Ordering an AWS Snowball device would not be a scalable or cost-effective solution for deploying the ML model. AWS Snowball is a service that provides physical devices for data transfer and edge computing, but it is not designed for continuous operation or frequent updates3.
* Deploying Amazon Monitron devices on each IP camera would not work because Amazon Monitron is a service that monitors the condition and performance of industrial equipment using sensors and machine learning, not cameras4.
References:
https://aws.amazon.com/greengrass/
https://docs.aws.amazon.com/greengrass/v2/developerguide/use-machine-learning-inference.html
https://aws.amazon.com/snowball/
https://aws.amazon.com/monitron/

**NEW QUESTION: 20**
Question:
A company runs an application on Amazon EC2 and AWS Lambda. The application stores temporary data in Amazon S3. The S3 objects are deleted after 24 hours.
The company deploys new versions of the application by launching AWS CloudFormation stacks. The stacks create the required resources. After validating a new version, the company deletes the old stack. The deletion of an old development stack recently failed.
A solutions architect needs to resolve this issue without major architecture changes.
Which solution will meet these requirements?
**A.** Create a Lambda function to delete objects from the S3 bucket. Add the Lambda function as a custom resource in the CloudFormation stack with a DependsOn attribute that points to the S3 bucket resource.

**B.** Modify the CloudFormation stack to attach a DeletionPolicy attribute with a value of Delete to the S3 bucket.
**C.** Update the CloudFormation stack to add a DeletionPolicy attribute with a value of Snapshot for the S3 bucket resource.
**D.** Update the CloudFormation template to create an Amazon EFS file system to store temporary files instead of Amazon S3. Configure the Lambda functions to run in the same VPC as the EFS file system.
**Answer: (SHOW ANSWER)**
Comprehensive and Detailed Explanation:
CloudFormation cannot delete non-empty S3 buckets. OptionAallows you to create acustom Lambda resourcethat deletes all objects in the S3 bucket before the stack deletes it. The DependsOn ensures the bucket deletion occurs only after the Lambda has completed.
* B: Adding DeletionPolicy: Delete does not resolve the issue if the bucket still contains objects.
* C: Snapshot doesn't apply to S3 and won't help here.
* D: Changing to Amazon EFS would require architectural changes, which are not allowed per requirements.
#Reference:
https://aws.amazon.com/blogs/devops/safely-delete-s3-buckets-using-aws-cloudformation/
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-deletionpolicy.html

**NEW QUESTION: 21**
A company's factory and automaton applications are running in a single VPC More than 23 applications run on a combination of Amazon EC2, Amazon Elastic Container Service (Amazon ECS), are Amazon RDS.
The company has software engineers spread across three teams. One of the three teams owns each application, and each team is responsible for the cost and performance of all of its applications. Team resources have tags that represent their application and team. The learns use IAH 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 item the last 12 months and to help forecast costs tor the next 12 months. A solution architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.
Which combination of actions will meet these requirement? Select THREE.)
**A.** Activate the user-defined 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.
**Answer: A,C,F (LEAVE A REPLY)**
To attribute AWS costs to specific applications or teams and enable detailed cost analysis and forecasting, the solution architect should recommend the following actions: A. Activating user-defined

cost allocation tags for resources associated with each application and team allows for detailed tracking of costs by these identifiers.
C: Creating a cost category for each application within AWS Billing and Cost Management enables the organization to group costs according to application, facilitating detailed reporting and analysis. F. Enabling Cost Explorer is essential for analyzing and visualizing AWS spending over time. It provides the capability to view historical costs and forecast future expenses, supporting the company's requirement for cost comparison and forecasting.
AWS Billing and Cost Management Documentation: Covers the activation of cost allocation tags, creation of cost categories, and the use of Cost Explorer for cost management.
AWS Tagging Strategies: Provides best practices for implementing tagging strategies that support cost allocation and reporting.
AWS Cost Explorer Documentation: Details how to use Cost Explorer to analyze and forecast AWS costs.

**NEW QUESTION: 22**
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 gatewayattachment. Which combination of steps will meet these requirements? (Select TWO)
**A.** From the management account, share the transit gateway with member accounts by using AWS Resource Access Manager
**B.** Prom 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 automatical^/ 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 automatical^ 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
**Answer: A,C (LEAVE A REPLY)**
https://aws.amazon.com/blogs/mt/self-service-vpcs-in-aws-control-tower-using-aws-service-catalog/
https://docs.aws.amazon.com/vpc/latest/tgw/tgw-transit-gateways.htmlhttps://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachment.html

**NEW QUESTION: 23**
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.
**Answer: A (LEAVE A REPLY)**
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 Deploying a NAT gateway and associating an Elastic IP address with it, and then configuring the VPC to use the NAT gateway, will give the application the ability to access the new service. This is because the NAT gateway will be the single public IP address that the external provider needs for the allow list. The NAT gateway will allow the application to access the service, while keeping the underlying Lambda functions private.
When configuring NAT gateways, you should ensure that the route table associated with the NAT gateway has a route to the internet gateway with a target of the internet gateway. Additionally, youshould ensure that the security group associated with the NAT gateway allows outbound traffic from the Lambda functions.
References:
AWS Certified Solutions Architect Professional Official Amazon Text Book[1], page 456
https://docs.aws.amazon.com/vpc/latest/userguide/VPC_NAT_Gateway.html

**NEW QUESTION: 24**
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 I business day, with loss of no more than I 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 tothe 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.
**Answer: (SHOW ANSWER)**
Using AWS Backup to create a scheduled daily backup plan for the EC2 instances will enable taking snapshots of the EC2 instances and their attached EBS volumes1. Configuring the backup task to copy the backups to a vault in the secondary Region will enable maintaining backups in a separate Region1. In the event of a failure, launching the CloudFormation template will enable deploying the network configuration in the secondary Region2. Restoring the instance volumes and configurations from the backup vault will enable recovering the EC2 instances and their data1. Transferring usage to the secondary Region will enable resuming operations2.

**NEW QUESTION: 25**
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.
**Answer: A (LEAVE A REPLY)**
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.

**NEW QUESTION: 26**
A solutions architect needs to review the design of an Amazon EMR cluster that is using the EMR File System (EMRFS). The cluster performs tasks that are critical to business needs. The cluster is running Amazon EC2 On-Demand Instances at all times tor all task, primary, and core nodes. The EMR tasks run each morning, starting at 1 ;00 AM. and take 6 hours to finish running. The amount of time to complete the processing is not a priority because the data is not referenced until late in the day.
The solutions architect must review the architecture and suggest a solution to minimize the compute costs.
Which solution should the solutions architect recommend to meet these requirements?
**A.** Launch all task, primary, and core nodes on Spool Instances in an instance fleet. Terminate the cluster, including all instances, when the processing is completed.
**B.** Launch the primary and core nodes on On-Demand Instances. Launch the task nodes on Spot Instances in an instance fleet. Terminate the cluster, including all instances, when the processing is completed.
Purchase Compute Savings Plans to cover the On-Demand Instance usage.
**C.** Continue to launch all nodes on On-Demand Instances. Terminate the cluster, including all instances, when the processing is completed. Purchase Compute Savings Plans to cover the On-Demand Instance usage
**D.** Launch the primary and core nodes on On-Demand Instances. Launch the task nodes on Spot Instances in an instance fleet. Terminate only the task node instances when the processing is completed.
Purchase Compute Savings Plans to cover the On-Demand Instance usage.
**Answer: (SHOW ANSWER)**
Amazon EC2 Spot Instances offer spare compute capacity at steep discounts compared to On-Demand prices.
Spot Instances can be interrupted by EC2 with two minutes of notification when EC2 needs the capacity back.
Amazon EMR can handle Spot interruptions gracefully by decommissioning the nodes and redistributing the tasks to other nodes. By launching all nodes on Spot Instances in an instance fleet, the solutions

architect can minimize the compute costs of the EMR cluster. An instance fleet is a collection of EC2 instances with different types and sizes that EMR automatically provisions to meet a defined target capacity. By terminating the cluster when the processing is completed, the solutions architect can avoid paying for idle resources.
References:
* https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-managed-scaling.html
* https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-instance-fleet.html
* https://aws.amazon.com/blogs/big-data/optimizing-amazon-emr-for-resilience-and-cost-with-capacity-optimized-spot-instances/

**NEW QUESTION: 27**
A company manufactures smart vehicles. The company uses a custom application to collect vehicle data. The vehicles use the MQTT protocol to connect to the application.
The company processes the data in 5-minute intervals. The company then copies vehicle telematics data to on- premises storage. Custom applications analyze this data to detect anomalies.
The number of vehicles that send data grows constantly. Newer vehicles generate high volumes of data. The on-premises storage solution is not able to scale for peak traffic, which results in data loss. The company must modernize the solution and migrate the solution to AWS to resolve the scaling challenges.
Which solution will meet these requirements with the LEAST operational overhead?
**A.** Use AWS IOT Greengrass to send the vehicle data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Create an Apache Kafka application to store the data in Amazon S3. Use a pretrained model in Amazon SageMaker to detect anomalies.
**B.** Use AWS IOT Core to receive the vehicle data. Configure rules to route data to an AmazonKinesis Data Firehose delivery stream that stores the data in Amazon S3. Create an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies.
**C.** Use AWS IOT FleetWise to collect the vehicle data. Send the data to an Amazon Kinesis data stream. Use an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use the built-in machine learning transforms in AWS Glue to detect anomalies.
**D.** Use Amazon MQ for RabbitMQ to collect the vehicle data. Send the data to an Amazon Kinesis Data Firehose delivery stream to store the data in Amazon S3. Use Amazon Lookout for Metrics to detect anomalies.
**Answer: B (LEAVE A REPLY)**
Using AWS IoT Core to receive the vehicle data will enable connecting the smart vehicles to the cloud using the MQTT protocol1. AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline2. Configuring rules to route data to an Amazon Kinesis Data Firehose delivery stream that stores the data in Amazon S3 will enable processing and storing the vehicle data in a scalable and reliable way3. Amazon Kinesis Data Firehose is a fully managed service that delivers real- time streaming data to destinations such as Amazon S3. Creating an Amazon Kinesis Data Analytics application that reads from the delivery stream to detect anomalies will enable analyzing the vehicle data using SQL queries or Apache Flink applications. Amazon Kinesis Data

Analytics is a fully managed service that enables you to process and analyze streaming data using SQL or Java.

**NEW QUESTION: 28**
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-contiguration 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.
**Answer: A (LEAVE A REPLY)**
https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with- alias-traffic-shifting/

**NEW QUESTION: 29**
A company has an organization in AWS Organizations that includes a separate AWS account for each of the company's departments. Application teams from different departments develop and deploy solutions independently.
The company wants to reduce compute costs and manage costs appropriately across departments. The company also wants to improve visibility into billing for individual departments. The company does not want to lose operational flexibility when the company selects compute resources.
Which solution will meet these requirements?
**A.** Configure AWS Organizations to use consolidated billing. Implement a tagging strategy that identifies departments. Use Tag Editor to apply tags to appropriate resources. Purchase Compute Savings Plans.
**B.** Use AWS Budgets for each department. Use SCPs to apply tags to appropriate resources. Purchase Compute Savings Plans.
**C.** Configure AWS Organizations to use consolidated billing. Implement a tagging strategy that identifies departments. Use SCPs to apply tags to appropriateresources. Purchase EC2 Instance Savings Plans.
**D.** Use AWS Budgets for each department. Use Tag Editor to apply tags to appropriate resources. Purchase EC2 Instance Savings Plans.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 30**
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 runs on Amazon EC2 instances running the Docker containers to process the data

**Answer: (SHOW ANSWER)**

AWS DataSync can be used to transfer the sequencing data to Amazon S3, which is a more efficient and faster method than using Snowball Edge devices. Once the data is in S3, S3 events can trigger an AWS Lambda function that starts an AWS Step Functions workflow. The Docker images can be stored in Amazon Elastic ContainerRegistry (Amazon ECR) and AWS Batch can be used to run the container and process the sequencing data.

**NEW QUESTION: 31**

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.

**Answer: B (LEAVE A REPLY)**

Access Analyzer is to assess the access policy.https://docs.aws.amazon.com/ja_jp/AmazonS3/latest/userguide/access-control-block-public-access.html



**NEW QUESTION: 32**

A company uses AWS CloudFormation to deploy applications within multiple VPCs that are all attached to a transit gateway Each VPC that sends traffic to the public internet must send the traffic through a shared services VPC Each subnet within a VPC uses the default VPC route table and the traffic is routed to the transit gateway The transit gateway uses its default route table for any VPC attachment A security audit reveals that an Amazon EC2 instance that is deployed within a VPC can communicate with an EC2 instance that is deployed in any of the company's other VPCs A solutions architect needs to limit the traffic between the VPCs. Each VPC must be able to communicate only with a predefined, limited set of authorized VPCs.

What should the solutions architect do to meet these requirements'?

**A.** Update the network ACL of each subnet within a VPC to allow outbound traffic only to the authorized VPCs Remove all deny rules except the default deny rule.

**B.** Update all the security groups that are used within a VPC to deny outbound traffic to security groups that are used within the unauthorized VPCs

**C.** Create a dedicated transit gateway route table for each VPC attachment. Route traffic only to the authorized VPCs.

**D.** Update the mam route table of each VPC to route traffic only to the authorized VPCs through the transit gateway

**Answer: C (LEAVE A REPLY)**

You can segment your network by creating multiple route tables in an AWS Transit Gateway and associate Amazon VPCs and VPNs to them. This will allow you to create isolated networks inside an AWS Transit Gateway similar to virtual routing and forwarding (VRFs) in traditional networks. The AWS Transit Gateway will have a default route table. The use of multiple route tables is optional.

**NEW QUESTION: 33**

A company provides a centralized Amazon EC2 application hosted in a single shared VPC The centralized application must be accessible from client applications running in the VPCs of other business units The centralized application front end is configured with a Network Load Balancer (NLB) for scalability Up to 10 business unit VPCs will need to be connected to the shared VPC Some ot the business unit VPC CIDR blocks overlap with the shared VPC and some overlap with each other Network connectivity to the centralized application in the shared VPC should be allowed from authorized business unit VPCs only Which network configuration should a solutions architect use to provide connectivity from the client applications in the business unit VPCs to the centralized application in the shared VPC?

**A.** Create an AWS Transit Gateway Attach the shared VPC and the authorized business unit VPCs to the transit gateway Create a single transit gateway route table and associate it with all of the attached VPCs Allow automatic propagation of routes from the attachments into the route table Configure VPC routing tables to send traffic to the transit gateway

**B.** Create a VPC endpoint service using the centralized application NLB and enable the option to require endpoint acceptance Create a VPC endpoint in each of the business unit VPCs using the service name of the endpoint service. Accept authorized endpoint requests from the endpoint serviceconsole.

**C.** Create a VPC peering connection from each business unit VPC to the shared VPC Accept the VPC peering connections from the shared VPC console Configure VPC routing tables to send traffic to the VPC peering connection

**D.** Configure a virtual private gateway for the shared VPC and create customer gateways for each of the authorized business unit VPCs Establish a Site-to-Site VPN connection from the business unit VPCs to the shared VPC Configure VPC routing tables to send traffic to the VPN connection

**Answer: (SHOW ANSWER)**

* Create VPC Endpoint Service:
* In the shared VPC, create a VPC endpoint service using the Network Load Balancer (NLB) that fronts the centralized application.
* Enable the option to require endpoint acceptance to control which business unit VPCs can connect to the service.
* Set Up VPC Endpoints in Business Unit VPCs:
* In each business unit VPC, create a VPC endpoint that points to the VPC endpoint service created in the shared VPC.
* Use the service name of the endpoint service created in the shared VPC for configuration.
* Accept Endpoint Requests:
* From the VPC endpoint service console in the shared VPC, review and accept endpoint connection requests from authorized business unit VPCs. This ensures that only authorized VPCs can access the centralized application.
* Configure Routing:
* Update the route tables in each business unit VPC to direct traffic destined for the centralized application through the VPC endpoint.

This solution ensures secure, private connectivity between the business unit VPCs and the shared VPC, even if there are overlapping CIDR blocks. It leverages AWS PrivateLink and VPC endpoints to provide scalable and controlled access(AWS Documentation)(Amazon Web Services, Inc.).

**NEW QUESTION: 34**
A financial services company sells its software-as-a-service (SaaS) platform for application compliance to large global banks. The SaaS platform runs on AWS and uses multiple AWS accounts that are managed in an organization in AWS Organizations. The SaaS platform uses many AWS resources globally.
For regulatory compliance, all API calls to AWS resources must be audited, tracked for changes, and stored in a durable and secure data store.
Which solution will meet these requirements with the LEAST operational overhead?
**A.** Create a new AWS CloudTrail trail. Use an existing Amazon S3 bucket in theorganization's management account to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 bucket.
**B.** Create a new AWS CloudTrail trail in each member account of the organization. Create new Amazon S3 buckets to store the logs. Deploy the trail to all AWS Regions. Enable MFA delete and encryption on the S3 buckets.
**C.** Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket with versioning turned on to store the logs. Deploy the trail for all accounts in the organization. Enable MFA delete and encryption on the S3 bucket.
**D.** Create a new AWS CloudTrail trail in the organization's management account. Create a new Amazon S3 bucket to store the logs. Configure Amazon Simple Notification Service (Amazon SNS) to send log-file delivery notifications to an external management system that will track the logs. Enable MFA delete and encryption on the S3 bucket.
**Answer: C (LEAVE A REPLY)**
The correct answer is C. This option uses AWS CloudTrail to create a trail in the organization's management account that applies to all accounts in the organization. This way, the company can centrally manage and audit all API calls to AWS resources across multiple accounts and regions. The company also needs to create a new Amazon S3 bucket with versioning turned on to store the logs. Versioning helps protect against accidental or malicious deletion of log files by keeping multiple versions of each object in the bucket. The company also needs to enable MFA delete and encryption on the S3 bucket to further enhance the security and durability of the data store.
Option A is incorrect because it uses an existing S3 bucket in the organization's management account to store the logs. This may not be optimal for regulatory compliance, as the existing bucket may have different permissions, encryption settings, or lifecycle policies than a dedicated bucket for CloudTrail logs.
Option B is incorrect because it requires creating a new CloudTrail trail in each member account of the organization. This adds operational overhead and complexity, as the company would need to manage multiple trails and S3 buckets across multiple accounts and regions.
Option D is incorrect because it requires configuring Amazon SNS to send log-file delivery notifications to an external management system that will track the logs. This adds unnecessary complexity and cost, as CloudTrail already provides log-file integrity validation and log-file digest delivery features that can help verify the authenticity and integrity of log files.
Reference: Creating a Trail for an Organization

**NEW QUESTION: 35**
A company has automated the nightly retraining of its machine learning models by using AWS Step Functions. The workflow consists of multiple steps that use AWS Lambda Each step can fail for various reasons and any failure causes a failure of the overall workflow A review reveals that the retraining has failed multiple nights in a row without the company noticing the failure A solutions architect needs to improve the workflow so that notifications are sent for all types of failures in the retraining process Which combination of steps should the solutions architect take to meet these requirements? (Select THREE)
**A.** Create an Amazon Simple Notification Service (Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list.
**B.** Create a task named "Email" that forwards the input arguments to the SNS topic
**C.** Add a Catch field all Task Map. and Parallel states that have a statement of "Error Equals": [ "States. ALL"] and "Next": "Email".
**D.** Add a new email address to Amazon Simple Email Service (Amazon SES). Verify the email address.
**E.** Create a task named "Email" that forwards the input arguments to the SES email address
**F.** Add a Catch field to all Task Map, and Parallel states that have a statement of "Error Equals": [ "states.Runtime"] and "Next": "Email".
**Answer: A,B,C (LEAVE A REPLY)**
* Create an Amazon Simple Notification Service (Amazon SNS) topic with a subscription of type "Email" that targets the team's mailing list. This will create a topic for sending notifications and add a subscription for the team's email list to that topic. C. Add a Catch field to all Task, Map, and Parallel states that have a statement of "ErrorEquals": [ "States.ALL" ] and "Next": "Email". This will ensure that any errors that occur in any of the steps in the workflow will trigger the "Email" task, which will forward the input arguments to the SNS topic created in step A. B. Create a task named "Email" that forwards the input arguments to the SNS topic. This will allow the company to send email notifications to the team's mailing list in case of any errors occurred in any step in the workflow.

**NEW QUESTION: 36**
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.
**Answer: B (LEAVE A REPLY)**
https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.htmlhttps://docs.aws.amazon.com/drs/latest/userguide/recovery-workflow-gs.html

**NEW QUESTION: 37**
A company wants to use Amazon Workspaces in combination with thin client devices to replace aging desktops. Employees use the desktops to access applications that work with clinical trial data. Corporate security policy states that access to the applications must be restricted to only company branch office locations. The company is considering adding an additional branch office in the next 6 months.
Which solution meets these requirements with the MOST operational efficiency?
**A.** Create an IP access control group rule with the list of public addresses from the branch offices. Associate the IP access control group with the Workspaces directory.
**B.** Use AWS Firewall Manager to create a web ACL rule with an IPSet with the list to public addresses from the branch office Locations-Associate the web ACL with the Workspaces directory.
**C.** Use AWS Certificate Manager (ACM) to issue trusted device certificates to the machines deployed in the branch office locations. Enable restricted access on the Workspaces directory.
**D.** Create a custom Workspace image with Windows Firewall configured to restrict access to the public addresses of the branch offices. Use the image to deploy the Workspaces.
**Answer: A (LEAVE A REPLY)**
Utilizing an IP access control group rule with the list of public addresses from branch offices and associating it with the Amazon WorkSpaces directory is the most operationally efficient solution. This method ensures that access to WorkSpaces is restricted to specified locations, aligning with the corporate security policy. This approach offers simplicity and flexibility, especially with the potential addition of a new branch office, as updating the IP access control group is straightforward.
AWS Documentation on Amazon WorkSpaces and IP Access Control Groups provides detailed instructions on how to implement access restrictions based on IP addresses. This solution aligns with best practices for securing virtual desktops while maintaining operational efficiency.

**NEW QUESTION: 38**
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? (Select 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
**Answer: (SHOW ANSWER)**
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.

**NEW QUESTION: 39**
A company has a complex web application that leverages Amazon CloudFront for global scalability and performance Over time, users report that the web application is slowing down The company's operations team reports that the CloudFront cache hit ratio has been dropping steadily. The cache metrics report indicates that query strings on some URLs are inconsistently ordered and are specified sometimes in mixed-case letters and sometimes in lowercase letters.
Which set of actions should the solutions architect take to increase the cache hit ratio as quickly as possible?
**A.** Deploy a Lambda@Edge function to sort parameters by name and force them lo be lowercase Select the CloudFront viewer request trigger to invoke the function
**B.** Update the CloudFront distribution to disable caching based on query string parameters.
**C.** Deploy a reverse proxy after the load balancer to post-process the emitted URLs in the application to force the URL strings to be lowercase.
**D.** Update the CloudFront distribution to specify casing-insensitive query string processing.
**Answer: A (LEAVE A REPLY)**
because Amazon CloudFront considers the case of parameter names and values when caching based on query string parameters , thus inconsistent query strings may cause CloudFront to forward mixed-cased/misordered requests to the origin. Triggering a Lambda@Edge function based on a viewer request

event to sort parameters by name and force them to be lowercase is the best choice.https://docs.aws.amazon.com /AmazonCloudFront/latest/DeveloperGuide/QueryStringParameters.html#query-string-parameters-optimizing- cachinghttps://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-cloudfront-trigger- events.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html#lambda- examples-normalize-query-string-parameters

**NEW QUESTION: 40**
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 MOTT 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 MOTT 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 MOTT 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.
**Answer: A (LEAVE A REPLY)**
it describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.
Reference:
https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/
https://aws.amazon.com/autoscaling/
https://aws.amazon.com/route53/

**NEW QUESTION: 41**
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 (Amazon CloudWatch Events) 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 accountsAssume the rotes 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
**Answer: (SHOW ANSWER)**
https://aws.amazon.com/solutions/implementations/automations-for-aws-firewall-manager/ 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

**NEW QUESTION: 42**
A company is running a workload that consists of thousands of Amazon EC2 instances. The workload is running in a VPC that contains several public subnets and private subnets. The public subnets have a route for

0.0.0.0/0 to an existing internet gateway. The private subnets have a route for 0.0.0.0/0 to an existing NAT gateway.
A solutions architect needs to migrate the entire fleet of EC2 instances to use IPv6. The EC2 instances that are in private subnets must not be accessible from the public internet.
What should the solutions architect do to meet these requirements?
**A.** Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets.Create a new NAT gateway, and enable IPv6 support. Update the VPC route tables for all private subnets, and add a route for ::/0 to the IPv6-enabled NAT gateway.
**B.** Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Create an egress-only internet gateway. Update the VPC route tables for all private subnets, and add a route for ::/0 to the egress-only internet gateway.
**C.** Update the existing VPC, and associate a custom IPv6 CIDR block with the VPC and all subnets. Update all the VPC route tables, and add a route for ::/0 to the internet gateway.
**D.** Update the existing VPC, and associate an Amazon-provided IPv6 CIDR block with the VPC and all subnets. Update the VPC route tables for all private subnets, and add a route for ::/0 to the NAT gateway.
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 43**
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.
**Answer: (SHOW ANSWER)**
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.

**NEW QUESTION: 44**
A company has AWS accounts that are in an organization in AWS Organizations. The company wants to track Amazon EC2 usage as a metric. The company's architecture team must receive a daily alert if the EC2 usage is more than 10% higher than the average EC2 usage from the last 30 days.
Which solution will meet these requirements?
**A.** Configure AWS Budgets in the organization's management account. Specify a usage type of EC2 running hours. Specify a daily period. Set the budget amountto be 10% more than the reported average usage for the last 30 days from AWS Cost Explorer. Configure an alert to notify the architecture team if the usagethreshold is met.
**B.** Configure AWS Cost Anomaly Detection in the organization's management account. Configure a monitor type of AWS Service. Apply a filter of Amazon EC2.Configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days.
**C.** Enable AWS Trusted Advisor in the organization's management account. Configure a cost optimization advisory alert to notify the architecture team if the EC2usage is 10% more than the reported average usage for the last 30 days.
**D.** Configure Amazon Detective in the organization's management account. Configure an EC2 usage anomaly alert to notify the architecture team ifDetectiveidentifies a usage anomaly of more than 10%.
**Answer: (SHOW ANSWER)**
AWS Cost Anomaly Detection is a feature of the AWS Cost Management suite that leverages machine learning to enable continuous monitoring of your AWS costs and usage, allowing you to identify unexpected and abnormal spending1. You can create cost monitors that evaluate specific AWS services, member accounts, cost allocation tags, orcost categories based on your AWS account structure2. You can also configure alert subscriptions that notify you when a cost monitor detects an anomaly that meets your threshold2. In this case, you can create a cost monitor with a monitor type of AWS Service and apply a filter of Amazon EC2 to track the EC2 usage as a metric. You can then configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days, which is the anomaly detection period used by AWS Cost Anomaly Detection3.

**NEW QUESTION: 45**
A company runs an loT platform on AWS loT 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 showhigh 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? {Select 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
**Answer: B,E (LEAVE A REPLY)**
* Option C is correct because leveraging Amazon Kinesis Data Streams and AWS Lambda to ingest and process the raw data resolves the issues permanently and enable growth as new sensors are provisioned.
Amazon Kinesis Data Streams is a serverless streaming data service that simplifies the capture, processing, and storage of data streams at any scale. Kinesis Data Streams can handle any amount of streaming data and process data from hundreds of thousands of sources with very low latency. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. Lambda can be triggered by Kinesis Data Streams events and process the data records in real time. Lambda can also scale automatically based on the incoming data volume. By using Kinesis Data Streams and Lambda, the company can reduce the load on the API servers and improve the performance and scalability of the data ingestion and processing layer3
* Option E is correct because re-architecting the database tier to use Amazon DynamoDB instead of an RDS MySQL DB instance resolves the issues permanently and enable growth as new sensors are provisioned. Amazon DynamoDB is a fully managed key-value and document database that delivers single-digit millisecond performance at any scale. DynamoDB supports auto scaling, which automatically adjusts read and write capacity based on actual traffic patterns. DynamoDB also supports on-demand capacity mode, which instantly accommodates up to double the previous peak traffic on a table. By using DynamoDB instead of RDS MySQL DB instance, the company can eliminate high write latency and improve scalability and performance of the database tier.
References: 1: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html 2: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/CHAP_AuroraOverview.html 3: https://docs.aws.amazon.com/streams/latest/dev/introduction.html :
https://docs.aws.amazon.com/lambda
/latest/dg/welcome.html : https://docs.aws.amazon.com/xray/latest/devguide/aws-xray.html :
https://docs.aws.
amazon.com/amazondynamodb/latest/developerguide/Introduction.html :

**NEW QUESTION: 46**
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' performance for these uploads Which solutions will meet these requirements? (Select 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
**Answer: A,D (LEAVE A REPLY)**
https://aws.amazon.com/premiumsupport/knowledge-center/s3-upload-large-files/ 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 these uploads by leveraging Amazon CloudFront's globally distributed edge locations to accelerate the uploads. Breaking the video files into chunks and using a multipart upload to transfer files to Amazon S3 will also improve the app's performance by allowing parts of the file to be uploaded in parallel, reducing the overall upload time.



**NEW QUESTION: 47**
A research company is running daily simul-ations in the AWS Cloud to meet high demand. The simu-lations run on several hundred Amazon EC2 instances that are based on Amazon Linux 2. Occasionally, a simu- lation gets stuck and requires a cloud operations engineer to solve the problem by connecting to an EC2 instance through SSH.
Company policy states that no EC2 instance can use the same SSH key and that all connections must be logged in AWS CloudTrail.
How can a solutions architect meet these requirements?
**A.** Set up AWS Secrets Manager to store the EC2 SSH key. Create a new AWS Lambda function to create a new SSH key and to call AWS Systems Manager Session Manager to set the SSH key on the EC2 instance. Configure Secrets Manager to use the Lambda function for automatic rotation once daily.Instruct the engineers to fetch the SSH key from Secrets Manager when they connect through any SSH client.
**B.** Launch new EC2 instances without setting up any SSH key for the instances. Set up EC2 Instance Connect on each instance. Create a new IAM policy, and attach it to the engineers' IAM role with an Allow statement for the SendSSHPublicKey action. Instruct the engineers to connect to the instance by using a browser-based SSH client from the EC2 console.
**C.** Launch new EC2 instances, and generate an individual SSH key for each instance. Store the SSH key in AWS Secrets Manager. Create a new IAM policy, and attach it to the engineers' IAM role with an Allow

statement for the GetSecretValue action. Instruct the engineers to fetch the SSH key from Secrets Manager when they connect through any SSH client.
**D.** Create an AWS Systems Manager document to run commands on EC2 instances to set a new unique SSH key. Create a new IAM policy, and attach it to the engineers' IAM role with an Allow statement to run Systems Manager documents. Instruct the engineers to run the document to set an SSH key and to connect through any SSH client.
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 48**
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.
**Answer: C (LEAVE A REPLY)**
The company should configure Amazon FSx for Lustre with an import and export policy. The company should link the new file system to an S3 bucket. The company should install the Lustre client and mount the document store to an Amazon EC2 instance by using NFS. This solution will meet the requirements with the least amount of effort becauseAmazon FSx for Lustre is a fully managed service that provides a high- performance filesystem optimized for fast processing of workloads such as machine learning, highperformance computing, video processing, financial modeling, and electronic design automation1
. Amazon FSx for Lustre can be linked to an S3 bucket and can import data from and export data to the bucket
2. The import and export policy can be configured to automatically import new or changed objects from S3 and export new or changed files to S33. This will ensure that the files are available to the public for download within 30 minutes. Amazon FSx for Lustre supports NFS version 3.0 protocol for Linux clients.
The other options are not correct because:

* Migrating the application to an AWS Lambda function would require a lot of effort and may not be feasible for the existing server that generates many documents. Lambda functions have limitations on execution time, memory, disk space, and network bandwidth.
* Setting up an Amazon S3 File Gateway would not work because S3 File Gateway does not support write-back caching, which means that files written to the file share are uploaded to S3 immediately and are not available locally until they are downloaded again. This would not provide fast local access to the files that the server generates and modifies.
* Configuring AWS DataSync to connect to an Amazon EC2 instance would not meet the requirement of making the files available to the public for download within 30 minutes. DataSync is a service that transfers data between on-premises storage systems and AWS storage services over the internet or AWS Direct Connect. DataSync tasks can be scheduled to run at specific times or intervals, but they are not triggered by file changes.
References:
https://aws.amazon.com/fsx/lustre/
https://docs.aws.amazon.com/fsx/latest/LustreGuide/create-fs-linked-data-repo.html
https://docs.aws.amazon.com/fsx/latest/LustreGuide/import-export-data-repositories.html
https://docs.aws.amazon.com/fsx/latest/LustreGuide/mounting-on-premises.html
https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html
https://docs.aws.amazon.com/storagegateway/latest/userguide/StorageGatewayConcepts.html
https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html

**NEW QUESTION: 49**
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 1AM 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? (Select TWO.)
**A.** Turn on the cross-origin resource sharing (CORS) feature for the S3 bucket in Account
**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 Bt 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/*"
    ]
}
```

**Answer: C,D (LEAVE A REPLY)**

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

**NEW QUESTION: 50**

A company is deploying a new web-based application and needs a storage solution for the Linux application servers. The company wants to create a single location for updates to application data for all instances. The active dataset will be up to 100 GB in size. A solutions architect has determined that peak operations will occur for 3 hours daily and will require a total of 225 MiBps of read throughput.

The solutions architect must design a Multi-AZ solution that makes a copy of the data available in another AWS Region for disaster recovery (DR). The DR copy has an RPO of less than 1 hour.
Which solution will meet these requirements?
**A.** Deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. Configure the file system for 75 MiBps of provisioned throughput. Implement replication to a file system in the DR Region.
**B.** Deploy a new Amazon FSx for Lustre file system. Configure Bursting Throughput mode for the file system. Use AWS Backup to back up the file system to the DR Region.
**C.** Deploy a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput. Enable Multi-Attach for the EBS volume. Use AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region.
**D.** Deploy an Amazon FSx for OpenZFS file system in both the production Region and the DR Region.Create an AWS DataSync scheduled task to replicate thedata from the production file system to the DR file system every 10 minutes.
**Answer: A (LEAVE A REPLY)**
The company should deploy a new Amazon Elastic File System (Amazon EFS) Multi-AZ file system. The company should configure the file system for 75 MiBps of provisioned throughput. The company should implement replication to a file system in the DR Region. This solution will meet the requirements because Amazon EFS is a serverless, fully elastic file storage service that lets you share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files1. By deploying a new Amazon EFS Multi-AZ file system, the company can create a single location for updates to application data for all instances. A Multi-AZ file system replicates data across multiple Availability Zones (AZs) within a Region, providing high availability and durability2. By configuring the file system for 75 MiBps of provisioned throughput, the company can ensure that it meets the peak operations requirement of 225 MiBps of read throughput. Provisioned throughput is a feature that enables you to specify a level of throughput that the file system can drive independent of the file system's size or burst credit balance3. By implementing replication to a file system in the DR Region, the company can make a copy of the data available in another AWS Region for disaster recovery. Replication is a feature that enables you to replicate data from one EFS file system to another EFS file system across AWS Regions. The replication process has an RPO of less than
1 hour.
The other options are not correct because:
* Deploying a new Amazon FSx for Lustre file system would not provide a single location for updates to application data for all instances. Amazon FSx for Lustre is a fully managed service that provides cost-effective, high-performance storage for compute workloads. However, it does not support concurrent write access from multiple instances. Using AWS Backup to back up the file system to the DR Region would not provide real-time replication of data. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. However, it does not support continuous data replication or cross-Region disaster recovery.
* Deploying a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume with 225 MiBps of throughput would not provide a single location for updates to application data for all instances.

Amazon EBS is a service that provides persistent block storage volumes for use with Amazon EC2 instances. However, it does not support concurrent access from multiple instances, unless Multi-Attach is enabled. Enabling Multi-Attach for the EBS volume would not provide Multi-AZ resilience or cross-Region replication. Multi-Attach is a feature that enables you toattach an EBS volume to multiple EC2 instances within the same Availability Zone. Using AWS Elastic Disaster Recovery to replicate the EBS volume to the DR Region would not provide real-time replication of data. AWS Elastic Disaster Recovery (AWS DRS) is a service that enables you to orchestrate and automate disaster recovery workflows across AWS Regions. However, it does not support continuous data replication or sub-hour RPOs.
* Deploying an Amazon FSx for OpenZFS file system in both the production Region and the DR Region would not be as simple or cost-effective as using Amazon EFS. Amazon FSx for OpenZFS is a fully managed service that provides high-performance storage with strong data consistency and advanced data management features for Linux workloads. However, it requires more configuration and management than Amazon EFS, which is serverless and fully elastic. Creating an AWS DataSync scheduled task to replicate the data from the production file system to the DR file system every 10 minutes would not provide real-time replication of data. AWS DataSync is a service that enables you to transfer data between on-premises storage and AWS services, or between AWS services. However, it does not support continuous data replication or sub-minute RPOs.
References:
https://aws.amazon.com/efs/
https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-azs
https://docs.aws.amazon.com/efs/latest/ug/performance.html#provisioned-throughput
https://docs.aws.amazon.com/efs/latest/ug/replication.html
https://aws.amazon.com/fsx/lustre/
https://aws.amazon.com/backup/
https://aws.amazon.com/ebs/
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volumes-multi.html

**NEW QUESTION: 51**
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 DynamoD8 Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoD8 table.

**Answer: B (LEAVE A REPLY)**

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.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 becomeunavailable or lose cached data.

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html

**NEW QUESTION: 52**

A solutions architect is preparing to deploy a new security tool into several previously unused AWS Regions.

The solutions architect will deploy the tool by using an AWS CloudFormation stack set. The stack set's template contains an 1AM role that has a custom name. Upon creation of the stack set. no stack instances are created successfully.

What should the solutions architect do to deploy the stacks successfully?

**A.** Enable the new Regions in all relevant accounts. Specify the CAPABILITY_NAMED_IAM capability during the creation of the stack set.

**B.** Use the Service Quotas console to request a quota increase for the number of CloudFormation stacks in each new Region in all relevant accounts. Specify the CAPABILITYJAM capability during the creation of the stack set.

**C.** Specify the CAPABILITY_NAMED_IAM capability and the SELF_MANAGED permissions model during the creation of the stack set.

**D.** Specify an administration role ARN and the CAPABILITYJAM capability during the creation of the stack set.

**Answer: A (LEAVE A REPLY)**

The CAPABILITY_NAMED_IAM capability is required when creating or updating CloudFormation stacks that contain IAM resources with custom names. This capability acknowledges that the template might create IAM resources that have broad permissions or affect other resources in the AWS account. The stack set's template contains an IAM role that has a custom name, so this capability is needed. Enabling the new Regions in all relevant accounts is also necessary to deploy the stack set across multiple Regions and accounts.

Option B is incorrect because the Service Quotas console is used to view and manage the quotas for AWS services, not for CloudFormation stacks. The number of stacks per Region per account is not a service quota that can be increased.

Option C is incorrect because the SELF_MANAGED permissions model is used when the administrator wants to retain full permissions to manage stack sets and stack instances. This model does not affect the creation of the stack set or the requirement for the CAPABILITY_NAMED_IAM capability.

Option D is incorrect because an administration role ARN is optional when creating a stack set. It is used to specify a role that CloudFormation assumes to create stack instances in the target accounts. It does not affect the creation of the stack set or the requirement for the CAPABILITY_NAMED_IAM capability.
:
1: AWS CloudFormation stack sets
2: Acknowledging IAM resources in AWS CloudFormation templates
3: AWS CloudFormation stack set permissions

**NEW QUESTION: 53**
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 Single Sign-On (AWS SSO) 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 Single Sign-On (AWS SSO) by using AWS SSO 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 AWS SSO 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.
**Answer: D (LEAVE A REPLY)**
https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/

**NEW QUESTION: 54**
A company recently wanted a web application from an on-premises data center to the AWS Cloud. The web application infrastructure consists of an Amazon CloudFront distribution that routes to an Application Load Balancer (ALB), with Amazon Elastic Container Service (Amazon ECS) to process requests. A recent security audit revealed that the web application is accessible by using both CloudFront and ALB endpoints.

However. the company requires that the web application must be accessible only by using the CloudFront endpoint.
Which solution will meet this requirement with the LEAST amount of effort?
**A.** Create a new security group and attach it to the CloudFront distribution. Update the ALB security group ingress to allow access only from the CloudFront security group.
**B.** Update ALB security group ingress to allow access only from the CloudFront managed prefix list.
**C.** Create a VPC interface endpoint for Elastic Load Balancing. Update the ALB scheme from internet-facing to internal_
**D.** Extract CloudFront IPS from the AWS provided ip-ranges.json document. Update ALB security group ingress to allow access only from CloudFront IPs.
**Answer: (SHOW ANSWER)**
The CloudFront managed prefix list contains the IP ranges for all CloudFront edge locations. By updating the ALB security group ingress to allow access only from this prefix list, the web application will be accessible only by using the CloudFront endpoint. This solution requires the least amount of effort compared to the other options, which involve creating new resources or updating existing ones. This solution also avoids hard- coding IP addresses, which can change over time.
Reference: section "Security and Compliance"

**NEW QUESTION: 55**
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 DocumentD3 (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.
**Answer: C (LEAVE A REPLY)**
On-demand capacity mode is the function of Dynamodb.https://aws.amazon.com/blogs/news/running-spiky- workloads-and-optimizing-costs-by-more-than-90-using-amazon-dynamodb-on-demand-capacity-mode/ Amazon DocumentDB Elastic Clustershttps://aws.amazon.com/blogs/news/announcing-amazon-documentdb- elastic-clusters/ Deploy Amazon EC2 instances in an Auto Scaling group across multiple

Availability Zones for the Java backend application. This will provide high availability and scalability, while allowing the company to retain the same database structure as the original application.

**NEW QUESTION: 56**
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.
**Answer: C,E (LEAVE A REPLY)**
"Save your custom error pages in a location that is accessible to CloudFront. We recommend that you store them in an Amazon S3 bucket, and that you don't storethem in the same place as the rest of your website or application's content. If you store the custom error pages on the same origin as your website or application, and the origin starts to return 5xx errors, CloudFront can't get the custom error pages because the origin server is unavailable."https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html

**NEW QUESTION: 57**
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.
**Answer: (SHOW ANSWER)**
https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and-phased- deployments-with-aws-codedeploy/

**NEW QUESTION: 58**
A company is designing an AWS Organizations structure. The company wants to standardize a process to apply tags across the entire organization. The company will require tags with specific values when a user creates a new resource. Each of the company's OUs will have unique tag values.
Which solution will meet these requirements?
**A.** Use an SCP to deny the creation of resources that do not have the required tags. Create a tag policy that Includes the tag values that the company has assigned to each OU. Attach the tag policies to the OUs.
**B.** Use an SCP to deny the creation of resources that do not have the required tags. Create a tag policy that includes the tag values that the company has assigned to each OU. Attach the tag policies to the organization's management account.
**C.** Use an SCP to allow the creation of resources only when the resources have the required tags. Create a tag policy that includes the tag values that the company has assigned to each OU. Attach the tag policies to the OUs.
**D.** Use an SCP to deny the creation of resources that do not have the required tags. Define the list of tags.Attach the SCP to the OUs
**Answer: A (LEAVE A REPLY)**
https://aws.amazon.com/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and-service-control-policies-scps/

**NEW QUESTION: 59**

A utility company collects usage data from smart meters every 5 minutes. Data is sent to API Gateway, processed by Lambda, and stored in DynamoDB. As usage increased, Lambda durations increased and DynamoDB PUTs failed with ProvisionedThroughputExceededException. Lambda also experiences TooManyRequestsException errors.
Which combination of changes will resolve these issues? (Select TWO.)
**A.** Increase the write capacity units to the DynamoDB table.
**B.** Increase the memory available to the Lambda functions.
**C.** Increase the payload size from the smart meters.
**D.** Stream the data into an Amazon Kinesis data stream from API Gateway and process the data in batches.
**E.** Collect data in an Amazon SQS FIFO queue, which triggers a Lambda function to process each message.
**Answer: A,D (LEAVE A REPLY)**
Comprehensive and Detailed in Depth Explanation:
A is correct because write throttling on DynamoDB means WCU (write capacity units) are insufficient. Increasing them will reduce the error rate.
D is correct because introducing Kinesis allows for efficient, high-throughput ingestion and batch processing, reducing the number of Lambda invocations and overall load.
References:
DynamoDB Throttling
Using Kinesis with Lambda

**NEW QUESTION: 60**
A company is using Amazon API Gateway to deploy a private REST API that will provide access to sensitive data. The API must be accessible only from an application that is deployed in a VPC. The company deploys the API successfully. However, the API is not accessible from an Amazon EC2 instance that is deployed in the VPC.
Which solution will provide connectivity between the EC2 instance and the API?
**A.** Create an interface VPC endpoint for API Gateway. Attach an endpoint policy that allows apigateway:* actions. Disable private DNS naming for the VPC endpoint. Configure an API resource policy that allows access from the VPC. Use the VPC endpoint's DNS name to access the API.
**B.** Create an interface VPC endpoint for API Gateway. Attach an endpoint policy that allows the execute-api:Invoke action. Enable private DNS naming for the VPC endpoint. Configure an API resource policy that allows access from the VPC endpoint. Use the API endpoint's DNS names to access the API. Most Voted
**C.** Create a Network Load Balancer (NLB) and a VPC link. Configure private integration between API Gateway and the NLB. Use the API endpoint's DNS names to access the API.
**D.** Create an Application Load Balancer (ALB) and a VPC Link. Configure private integration between API Gateway and the ALB. Use the ALB endpoint's DNS name toaccess the API.
**Answer: (SHOW ANSWER)**

According to the AWS documentation1, to access a private API from a VPC, you need to do the following:
* Create an interface VPC endpoint for API Gateway in your VPC. This creates a private connection between your VPC and API Gateway.
* Attach an endpoint policy to the VPC endpoint that allows the execute-api:lnvoke action for your private API. This grants permission to invoke your API from the VPC.
* Enable private DNS naming for the VPC endpoint. This allows you to use the same DNS names for your private APIs as you would for public APIs.
* Configure a resource policy for your private API that allows access from the VPC endpoint. This controls who can access your API and under what conditions.
* Use the API endpoint's DNS names to access the API from your VPC. For example, https://api-id.execute-api.region.amazonaws.com/stage.

**NEW QUESTION: 61**
A financial services company runs a complex, multi-tier application on Amazon EC2 instances and AWS Lambda functions. The application stores temporary data in Amazon S3. The S3 objects are valid for only 45 minutes and are deleted after 24 hours.
The company deploys each version of the application by launching an AWS CloudFormation stack. The stack creates all resources that are required to run the application. When the company deploys and validates a new application version, the company deletes the CloudFormation stack of the old version.
The company recently tried to delete the CloudFormation stack of an old application version, but the operation failed. An analysis shows that CloudFormation failed to delete an existing S3 bucket. A solutions architect needs to resolve this issue without making major changes to the application's architecture.
Which solution meets these requirements?
**A.** Implement a Lambda function that deletes all files from a given S3 bucket. Integrate this Lambda function as a custom resource into the CloudFormation stack. Ensure that the custom resource has a DependsOn attribute that points to the S3 bucket's resource.
**B.** Modify the CloudFormation template to provision an Amazon Elastic File System (Amazon EFS) file system to store the temporary files there instead of in Amazon S3. Configure the Lambda functions to run in the same VPC as the file system. Mount the file system to the EC2 instances and Lambda functions.
**C.** Modify the CloudFormation stack to create an S3 Lifecycle rule that expires all objects 45 minutes after creation. Add a DependsOn attribute that points to the S3 bucket's resource.
**D.** Modify the CloudFormation stack to attach a DeletionPolicy attribute with a value of Delete to the S3 bucket.
**Answer: D (LEAVE A REPLY)**
Explanation: This option allows the solutions architect to use a DeletionPolicy attribute to specify how AWS CloudFormation handles the deletion of an S3 bucket when the stack is deleted1. By setting the value of Delete, the solutions architect can instruct CloudFormation to delete the bucket and all of its contents1. This option does not require any major changes to the application's architecture or any additional resources.

:
Deletion policies



**NEW QUESTION: 62**
A scientific company needs to process text and image data from an Amazon S3 bucket. The data is collected from several radar stations during a live, time-critical phase of a deep space mission. The radar stations upload the data to the source S3 bucket. The data is prefixed by radar station identification number.
The company created a destination S3 bucket in a second account. Data must be copied from the source S3 bucket to the destination S3 bucket to meet a compliance objective. The replication occurs through the use of an S3 replication rule to cover all objects in the source S3 bucket.
One specific radar station is identified as having the most accurate data. Data replication at this radar station must be monitored for completion within 30 minutes after the radar station uploads the objects to the source S3 bucket.
What should a solutions architect do to meet these requirements?
**A.** Set up an AWS DataSync agent to replicate the prefixed data from the source S3 bucket to the destination S3 bucket. Select to use all available bandwidth on the task, and monitor the task to ensure that it is in the TRANSFERRING status. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
**B.** In the second account, create another S3 bucket to receive data from the radar station with the most accurate data. Set up a new replication rule for this new S3 bucket to separate the replication from the other radar stations. Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.
**C.** Enable Amazon S3 Transfer Acceleration on the source S3 bucket, and configure the radar station with the most accurate data to use the new endpoint. Monitor theS3 destination bucket's TotalRequestLatency metric. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert if this status changes.
**D.** Create a new S3 replication rule on the source S3 bucket that filters for the keys that use the prefix of the radar station with the most accurate data. Enable S3 Replication Time Control (S3 RTC). Monitor the maximum replication time to the destination. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger an alert when the time exceeds the desired threshold.
**Answer: (SHOW ANSWER)**

https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication-time-control.html

**NEW QUESTION: 63**
A company's compliance audit reveals that some Amazon Elastic Block Store (Amazon EBS) volumes that were created in an AWS account were not encrypted. A solutions architect must Implement a solution to encrypt all new EBS volumes at rest Which solution will meet this requirement with the LEAST effort?
**A.** Create an Amazon EventBridge rule to detect the creation of unencrypted EBS volumes. Invoke an AWS Lambda function to delete noncompliant volumes.
**B.** Use AWS Audit Manager with data encryption.
**C.** Create an AWS Config rule to detect the creation of a new EBS volume. Encrypt the volume by using AWS Systems Manager Automation.
**D.** Turn in EBS encryption by default in all AWS Regions.
**Answer: D (LEAVE A REPLY)**
The most effortless way to ensure that all new Amazon Elastic Block Store (EBS) volumes are encrypted at rest is to enable EBS encryption by default in all AWS Regions. This setting automatically encrypts all new EBS volumes and snapshots created in the account, thereby ensuring compliance with encryption policies without the need for manual intervention or additional monitoring.
AWS Documentation on Amazon EBS encryption provides guidance on enabling EBS encryption by default. This approach aligns with AWS best practices for data protection and compliance, ensuring that all new EBS volumes adhere to encryption requirements with minimal operational effort.

**NEW QUESTION: 64**
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.

**Answer: D (LEAVE A REPLY)**

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.

**NEW QUESTION: 65**

A software development company has multiple engineers who ate working remotely. The company is running Active Directory Domain Services (AD DS) on an Amazon EC2 instance. The company's security policy states that al internal, nonpublic services that are deployed in a VPC must be accessible through a VPN. Multi- factor authentication (MFA) must be used for access to a VPN.

What should a solutions architect do to meet these requirements?

**A.** Create an AWS Sire-to-Site VPN connection. Configure Integration between a VPN and AD DS. Use an Amazon Workspaces client with MFA support enabled to establish a VPN connection.

**B.** Create an AWS Client VPN endpoint Create an AD Connector directory tor integration with AD DS. Enable MFA tor AD Connector. Use AWS Client VPN to establish a VPN connection.

**C.** Create multiple AWS Site-to-Site VPN connections by using AWS VPN CloudHub. Configure integration between AWS VPN CloudHub and AD DS. Use AWS Copilot to establish a VPN connection.

**D.** Create an Amazon WorkLink endpoint. Configure integration between Amazon WorkLink and AD DS.Enable MFA in Amazon WorkLink. Use AWS Client VPN to establish a VPN connection.

**Answer: B (LEAVE A REPLY)**

Setting up an AWS Client VPN endpoint and integrating it with Active Directory Domain Services (AD DS) using an AD Connector directory enables secure remote access to internal services deployed in a VPC. Enabling multi-factor authentication (MFA) for AD Connectorenhances security by adding an additional layer of authentication. This solution meets the company's requirements for secure remote access through a VPN with MFA, ensuring that the security policy is adhered to while providing a seamless experience for the remote engineers.

AWS Documentation on AWS Client VPN and AD Connector provides detailed instructions on setting up a Client VPN endpoint and integrating it with existing Active Directory for authentication. This solution aligns with AWS best practices for secure remote access to AWS resources.

**NEW QUESTION: 66**

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 CloudFormation template from an account that has the necessary IAM permissions
**D.** Use nested stacks with AWS CloudFormation templates Change the Region by using nestedstacks
**Answer: C (LEAVE A REPLY)**
https://aws.amazon.com/blogs/aws/new-use-aws-cloudformation-stacksets-for-multiple-accounts-in-an-aws- organization/ 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.

**NEW QUESTION: 67**
A company runs a serverless ecommerce application on AWS. The application uses API Gateway to invoke Java Lambda functions that connect to an Amazon RDS for MySQL database. During a sale event, traffic spikes caused slow performance and DB connection failures.
Which solution will improve performance with the LEAST application change?
**A.** Move DB connection outside Lambda handler and increase provisioned concurrency.
**B.** Use RDS Proxy. Store DB credentials in Secrets Manager. Update Lambda to use RDS Proxy. Increase provisioned concurrency.
**C.** Increase max_connections parameter in a custom DB parameter group and reboot. Increase reserved concurrency.
**D.** Use RDS Proxy and Secrets Manager. Increase reserved concurrency.
**Answer: B (LEAVE A REPLY)**
Comprehensive and Detailed in Depth Explanation:
B is correct because RDS Proxy allows for connection pooling and improved management of DB connections.
Lambda + RDS often runs into connection exhaustion during high concurrency unless RDS Proxy is used.
Provisioned concurrency prevents cold starts.
References:
RDS Proxy Overview
Provisioned Concurrency

**NEW QUESTION: 68**
A company recently started hosting new application workloads in the AWS Cloud. The company is using Amazon EC2 instances, Amazon Elastic File System (Amazon EFS) file systems, and Amazon RDS DB instances.

To meet regulatory and business requirements, the company must make the following changes for data backups:
* Backups must be retained based on custom daily, weekly, and monthly requirements.
* Backups must be replicated to at least one other AWS Region immediately after capture.
* The backup solution must provide a single source of backup status across the AWS environment.
* The backup solution must send immediate notifications upon failure of any resource backup.
Which combination of steps will meet this requirement with the LEAST amount of operational overhead? (Select THREE.)
**A.** Create an AWS Backup plan with a backup rule for each of the retention requirements.
**B.** Configure an AWS backup plan to copy backups to another Region.
**C.** Create an AWS Lambda function to replicate backups to another Region and send notification if a failure occurs.
**D.** Add an Amazon Simple Notification Service (Amazon SNS) topic to the backup plan to send a notification for finished jobs that have any status except BACKUP- JOB- COMPLETED.
**E.** Create an Amazon Data Lifecycle Manager (Amazon DLM) snapshot lifecycle policy for each of the retention requirements.
**F.** Set up RDS snapshots on each database.
**Answer: A,B,D (LEAVE A REPLY)**
Cross region with AWS Backup:https://docs.aws.amazon.com/aws-backup/latest/devguide/cross-region-backup.html

**NEW QUESTION: 69**
Question:
A company hosts an ecommerce site using EC2, ALB, and DynamoDB in one AWS Region. The site uses a custom domain in Route 53. The company wants toreplicate the stack to a second Regionfordisaster recoveryandfaster accessfor global customers.
What should the architect do?
**A.** Use CloudFormation to deploy to the second Region. Use Route 53 latency-based routing. Enable global tables in DynamoDB.
**B.** Use the console to recreate the infra manually in the second Region. Use weighted routing.
**C.** Replicate only the S3 and DynamoDB data. Use Route 53 failover routing.
**D.** Use Beanstalk and DynamoDB Streams for replication. Use latency-based routing.
**Answer: A (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
A is correct because:
* CloudFormation templatesenable repeatable infrastructure deployment.
* Route 53 latency-based routingensures users hit the closest Region.
* DynamoDB global tablesallow multi-Region, active-active replication of application data.
Manual console work (B) is not scalable.
C lacks EC2/ALB replication.
D adds unnecessary services like Beanstalk and doesn't scale cleanly.

#Reference:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html

**NEW QUESTION: 70**
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 CtoudWatch 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.
**Answer: (SHOW ANSWER)**
https://aws.amazon.com/es/blogs/industries/how-to-monitor-alert-and-remediate-non-compliant-hipaa-findings-on-aws/

**NEW QUESTION: 71**
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.** 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.
**C.** 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. Configurean S3 Multi-Region Access Point that uses both S3 buckets. Deploy a modified application to both Regions.

**D.** 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.
**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 72**
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 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? (Select TWO.)
**A.** Deploy the SOS 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 topics in each Region.
**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.
**Answer: (SHOW ANSWER)**
https://docs.aws.amazon.com/sns/latest/dg/sns-cross-region-delivery.html

**NEW QUESTION: 73**
A large payroll company recently merged with a small staffing company. The unified company now has multiple business units, each with its own existing AWS account.
A solutions architect must ensure that the company can centrally manage the billing and access policies for all the AWS accounts. The solutions architect configures AWS Organizations by sending an invitation to all member accounts of the company from a centralized management account.
What should the solutions architect do next to meet these requirements?
**A.** Create the OrganizationAccountAccess IAM group in each member account. Include the necessary IAM roles for each administrator.
**B.** Create the OrganizationAccountAccessRole IAM role in each member account. Grant permission to the management account to assume the IAM role.
**C.** Create the OrganizationAccountAccessRole IAM role in the management account. Attach the AdministratorAccess AWS managed policy to the IAM role.Assign the IAM role to the administrators in each member account.
**D.** Create the OrganizationAccountAccessPolicy IAM policy in each member account. Connect the member accounts to the management account by using cross-account access.
**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 74**
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.
**Answer: (SHOW ANSWER)**
The S3 Glacier Deep Archive storage class is the lowest-cost storage class offered by Amazon S3, and it is designed for archival data that is accessed infrequently and for which retrieval time of several hours is acceptable. S3 interface endpoint for the VPCensures that access to the bucket is only from resources within the VPC and this will meet the requirement of not being accessible to the public. And also, S3 bucket can be configured for website hosting, and this will allow employees to access the documents through the corporate intranet. Using an EC2 instance and a file system or block store would be more expensive and unnecessary because the number of requests to the data will be low and availability and speed of retrieval are not concerns.
Additionally, using Amazon S3 bucket will provide durability, scalability and availability of data.

**NEW QUESTION: 75**
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-hourperiod and is double the average load. The application load is close to the average load tor 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.
**Answer: D (LEAVE A REPLY)**
This solution meets the requirements by using Application Auto Scaling to automatically increase capacity during the peak period, which will handle the double the average load. And by purchasing reserved RCUs and WCUs to match the average load, it will minimize the cost of the table for the rest of the week when the load is close to the average.

**NEW QUESTION: 76**
A company has mounted sensors to collect information about environmental parameters such as humidity and light throughout all the company's factories. The company needs to stream and analyze the data in the AWS Cloud in real time. If any of the parameters fall out of acceptable ranges, the factory operations team must receive a notification immediately.
Which solution will meet these requirements?
**A.** Stream the data to an Amazon Kinesis Data Firehose delivery stream. Use AWS Step Functions to consume and analyze the data in the Kinesis Data Firehose delivery stream. use Amazon Simple Notification Service (Amazon SNS) to notify the operations team.
**B.** Stream the data to an Amazon Managed Streaming for Apache Kafka (Amazon MSK) cluster. Set up a trigger in Amazon MSK to invoke an AWS Fargate taskto analyze the data. Use Amazon Simple Email Service (Amazon SES) to notify the operations team.
**C.** Stream the data to an Amazon Kinesis data stream. Create an AWS Lambda function to consume the Kinesis data stream and to analyze the data. UseAmazon Simple Notification Service (Amazon SNS) to notify the operations team.
**D.** Stream the data to an Amazon Kinesis Data Analytics application. I-Jse an automatically scaled and containerized service in Amazon Elastic ContainerService (Amazon ECS) to consume and analyze the data. use Amazon Simple Email Service (Amazon SES) to notify the operations team.
**Answer: C (LEAVE A REPLY)**
The best solution is to stream the data to an Amazon Kinesis data stream and create an AWS Lambda function to consume the Kinesis data stream and to analyze the data. Amazon Kinesis is a web service that can collect, process, and analyze real-time streaming data from various sources, such as sensors. AWS Lambda is a serverless computing service that can run code in response to events, such as incoming data from a Kinesis data stream. By using AWS Lambda, the company can avoid provisioning or managing servers and scale automatically based on the demand. Amazon Simple Notification Service (Amazon SNS) is a web service that enables applications to send and receive notifications from the cloud. By using Amazon SNS, the company can notify the operations team immediately if any of the parameters fall out of acceptable ranges. This solution meets all the requirements of the company.
References: Amazon Kinesis Documentation, AWS Lambda Documentation, Amazon Simple Notification Service Documentation

**NEW QUESTION: 77**
An e-commerce company is revamping its IT infrastructure and is planning to use AWS services. The company's CIO has asked a solutions architect to design a simple, highly available, and loosely coupled order processing application. The application is responsible for receiving and processing orders before storing them in an Amazon DynamoDB table. The application has a sporadic traffic pattern and should be able to scale during marketing campaigns to process the orders with minimal delays.
Which of the following is the MOST reliable approach to meet the requirements?
**A.** Receive the orders in an Amazon EC2-hosted database and use EC2 instances to process them.
**B.** Receive the orders in an Amazon SQS queue and invoke an AWS Lambda function to process them.
**C.** Receive the orders using the AWS Step Functions program and launch an Amazon ECS container to process them.
**D.** Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.
**Answer: B (LEAVE A REPLY)**
The best option is to use Amazon SQS and AWS Lambda to create a serverless order processing application.
Amazon SQS is a fully managed message queue service that can decouple the order receiving and processing components, making the application more scalable and fault-tolerant. AWS Lambda is a serverless compute service that can automatically scale to handle the incoming messages from the SQS queue and process them according to the business logic. AWS Lambda can also integrate with Amazon DynamoDB to store the processed orders in a fast and flexible NoSQL database. Thisapproach eliminates the need to provision, manage, or scale any servers or containers, and reduces the operational overhead and cost.
Option A is not reliable because using an EC2-hosted database to receive the orders introduces a single point of failure and a scalability bottleneck. EC2 instances also require more management and configuration than serverless services.
Option C is not reliable because using AWS Step Functions to receive the orders adds unnecessary complexity and cost to the application. AWS Step Functions is a service that coordinates multiple AWS services into a serverless workflow, but it is not designed to handle high-volume, sporadic, or unpredictable traffic patterns. AWS Step Functions also charges per state transition, which can be expensive for a large number of orders. Launching an ECS container to process each order also requires more resources and management than invoking a Lambda function.

Option D is not reliable because using Amazon Kinesis Data Streams to receive the orders is not suitable for this use case. Amazon Kinesis Data Streams is a service that enables real-time processing of streaming data at scale, but it is not meant for asynchronous message queuing. Amazon Kinesis Data Streams requires consumers to poll the data from the stream, which can introduce latency and complexity. Amazon Kinesis Data Streams also charges per shard hour, which can be expensive for a sporadic traffic pattern.
:
Amazon SQS
AWS Lambda
Amazon DynamoDB
AWS Step Functions
Amazon ECS

**NEW QUESTION: 78**
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 Virealization 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(CIFSJ 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
**Answer: D (LEAVE A REPLY)**
https://docs.aws.amazon.com/vm-import/latest/userguide/vmimport-image-import.html
- Export an OVF Template
- Create / use an Amazon S3 bucket for storing the exported images. The bucket must be in the Region where you want to import your VMs.
- Create an IAM role named vmimport.
- You'll use AWS CLI to run the import commands.
https://aws.amazon.com/premiumsupport/knowledge-center/import-instances/

**NEW QUESTION: 79**

A company has a web application that uses Amazon API Gateway. AWS Lambda and Amazon DynamoDB A recent marketing campaign has increased demand Monitoring software reports that many requests have significantly longer response times than before the marketing campaign A solutions architect enabled Amazon CloudWatch Logs for API Gateway and noticed that errors are occurring on 20% of the requests. In CloudWatch. the Lambda function. Throttles metric represents 1% of the requests and the Errors metric represents 10% of the requests Application logs indicate that, when errors occur there is a call to DynamoDB What change should the solutions architect make to improve the current response times as the web application becomes more popular'?

**A.** Increase the concurrency limit of the Lambda function

**B.** Implement DynamoDB auto scaling on the table

**C.** Increase the API Gateway throttle limit

**D.** Re-create the DynamoDB table with a better-partitioned primary index.

**Answer: B (LEAVE A REPLY)**

* Enable DynamoDB Auto Scaling:
* Navigate to the DynamoDB console and select the table experiencing high demand.
* Go to the "Capacity" tab and enable auto scaling for both read and write capacity units. Auto scaling adjusts the provisioned throughput capacity automatically in response to actual traffic patterns, ensuring the table can handle the increased load.
* Configure Auto Scaling Policies:
* Set the minimum and maximum capacity units to define the range within which auto scaling can adjust the provisioned throughput.
* Specify target utilization percentages for read and write operations, typically around 70%, to maintain a balance between performance and cost.
* Monitor and Adjust:
* Use Amazon CloudWatch to monitor the auto scaling activity and ensure it is effectively handling the increased demand.
* Adjust the auto scaling settings if necessary to better match the traffic patterns and application requirements.

By enabling DynamoDB auto scaling, you ensure that the database can handle the fluctuating traffic volumes without manual intervention, improving response times and reducing errors.

References

* AWS Compute Blog on Using API Gateway as a Proxy for DynamoDB#60#.
* AWS Database Blog on DynamoDB Accelerator (DAX)#59#.

**NEW QUESTION: 80**

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.
**Answer: C (LEAVE A REPLY)**
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/what-is.htmlhttps://docs.aws.amazon.com /directoryservice/latest/admin-guide/ms_ad_join_instance.html

**NEW QUESTION: 81**
Question:
A company is modernizing a legacy.NET Frameworkapplication backed by SQL Server. Requirements:
* Containerize into microservices.
* Control OS patches and storage.
* Add load balancing.
* Ensure high availability.Which solution meets all of these with minimal refactoring?
**A.** Use App2Container to deploy on ECS EC2 with ALB and RDS for SQL Server.
**B.** Use App2Container on ECS EC2 with NLB and Aurora MySQL.
**C.** Use Porting Assistant and EKS with Fargate and Aurora MySQL.
**D.** Use Porting Assistant and EKS with Fargate and RDS SQL Server.
**Answer: A (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
A is correct because:
* App2Containersupports packaging .NET Framework apps into containers without porting to .NET Core.
* ECS with EC2gives full control over the OS and patching.

* ALBhandles microservice-level load balancing.
* RDS for SQL Serverwith Multi-AZ ensures high availability.
Options B, C, and D involve Aurora MySQL (incompatible with SQL Server features) orrequire .NET Core, which involves more significant application changes.
References:
App2Container Overview
ECS EC2 vs Fargate

**NEW QUESTION: 82**
A company is developing a gene reporting device that will collect genomic information to assist researchers with collecting large samples of data from a diverse population. The device will push 8 KB of genomic data every second to a data platform that will need to process and analyze the data and provide information back to researchers. The data platform must meet the following requirements:
*Provide near-real-time analytics of the inbound genomic data
*Ensure the data is flexible, parallel, and durable
*Deliver results of processing to a data warehouse
Which strategy should a solutions architect use to meet these requirements?
**A.** Use Amazon Kinesis Data Firehose to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon RDS instance.
**B.** Use Amazon Kinesis Data Streams to collect the inbound sensor data, analyze the data with Kinesis clients, and save the results to an Amazon Redshift cluster using Amazon EMR.
**C.** Use Amazon S3 to collect the inbound device data, analyze the data from Amazon SOS with Kinesis, and save the results to an Amazon Redshift cluster.
**D.** Use an Amazon API Gateway to put requests into an Amazon SQS queue, analyze the data with an AWS Lambda function, and save the results to an Amazon Redshift cluster using Amazon EMR.
**Answer: (SHOW ANSWER)**
Kinesis Data Streams is a real-time streaming service and provide near-real-time analytics. Also the question
"Deliver results of processing to a data warehouse" and this option has redshift cluster which is a powerful data warehousing solution that can handle large-scale analytics workloads.

**NEW QUESTION: 83**
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 fi-nance 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 con-trol). Establish trust with the marketing team's account. In the mar-keting team's account, create an IAM role that has permissions to as-sume the IAM role in the finance team's account.
**C.** Create a resource-based IAM policy that includes conditions for spe-cific 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 Dyna-moDB 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.
**Answer: C (LEAVE A REPLY)**
The company should create a resource-based IAM policy that includes conditions for specific DynamoDB attributes (fine-grained access control). The company should attach the policy to the DynamoDB table. In the marketing team's account, the company should create an IAM role that has permissions to access the DynamoDB table in the finance team's account. This solution will meet the requirements because a resource- based IAM policy is a policy that you attach to an AWS resource (such as a DynamoDB table) to control who can access that resource and what actions they can perform on it. You can use IAM policy conditions to specify fine-grained access control for DynamoDB items and attributes. For example, you can allow or deny access to specific attributes of all items in a table by matching on attribute names1. By creating a resource- based policy that allows access to only specific attributes of the DynamoDB table and attaching it to the table, the company can restrict access to confidential data. By creating an IAM role in the marketing team's account that has permissions to access the DynamoDB table in the finance team's account, the company can enable cross-account access.
The other options are not correct because:
* Creating an SCP to grant the marketing team's AWS account access to the specific attributes of the DynamoDB table would not work because SCPs are policies that you can use with AWS Organizations to manage permissions in your organization's accounts. SCPs do not grant permissions; instead, they specify the maximum permissions that identities in an account can have2. SCPs cannot be used to specify fine-grained access control for DynamoDB items and attributes.
* Creating an IAM role in the finance team's account by using IAM policy conditions for specific DynamoDB attributes and establishing trust with the marketing team's account would not work because IAM roles are identities that you can create in your account that have specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources3. However, creating an IAM role in the finance team's account would not restrict access to specific attributes of the DynamoDB table; it would only allow cross-account access. The company would still need a resource-based policy attached to the table to enforce fine-grained access control.
* Creating an IAM role in the finance team's account to access the DynamoDB table and using an IAM permissions boundary to limit the access to the specific attributes would not work because IAM permissions boundaries are policies that you use to delegate permissions management to other users.

You can use permissions boundaries to limit the maximum permissions that an identity-based policy can grant to an IAM entity (user or role)4. Permissions boundaries cannot be used to specify fine- grained access control for DynamoDB items and attributes.
References:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/specifying-conditions.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

**NEW QUESTION: 84**
A company is running an application that uses an Amazon ElastiCache for Redis cluster as acaching 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
**Answer: B (LEAVE A REPLY)**
Creating an AUTH token and storing it in AWS Secrets Manager and configuring the existing cluster to use the AUTH token and configure encryption in transit, and updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, would meet the requirements for user authentication and end-to-end encryption.
AWS Secrets Manager is a service that enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Secrets Manager also enables you to encrypt the data and ensure that only authorized users and applications can access it.
By configuring the existing cluster to use the AUTH token and encryption in transit, all data will be encrypted as it is sent over the network, providing additional security for the data stored in ElastiCache.

Additionally, by updating the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication, it ensures that only authorized users and applications can access the cache.
Reference:
AWS Secrets Manager documentation:https://aws.amazon.com/secrets-manager/ Encryption in transit for ElastiCache:https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug
/encryption.html
Authentication and Authorization for
ElastiCache:https://docs.aws.amazon.com/AmazonElastiCache/latest
/red-ug/accessing-elasticache.html

**NEW QUESTION: 85**
A company uses AWS Organizations to manage a multi-account structure. The company has hundreds of AWS accounts and expects the number of accounts to increase. The company is building a new application that uses Docker images. The company will push the Docker images to Amazon Elastic Container Registry (Amazon ECR). Only accounts that are within the company's organization should have access to the images.
The company has a CI/CD process that runs frequently. The company wants to retain all the tagged images.
However, the company wants to retain only the five most recent untagged images.
Which solution will meet these requirements with the LEAST operational overhead?
**A.** Create a private repository in Amazon ECR. Create a permissions policy for the repository that allows only required ECR operations. Include a condition to allow the ECR operations if the value of the aws: PrincipalOrgID condition key is equal to the ID of the company's organization. Add a lifecycle rule to the ECR repository that deletes all untagged images over the count of five.
**B.** Create a public repository in Amazon ECR. Create an IAM role in the ECR account. Set permissions so that any account can assume the role if the value of the aws:PrincipalOrgID condition key is equal to the ID of the company's organization. Add a lifecycle rule to the ECR repository that deletes all untagged images over the count of five.
**C.** Create a private repository in Amazon ECR. Create a permissions policy for the repository that includes only required ECR operations. Include a condition to allow the ECR operations for all account IDs in the organization. Schedule a daily Amazon EventBridge rule to invoke an AWS Lambda function that deletes all untagged images over the count of five.
**D.** Create a public repository in Amazon ECR. Configure Amazon ECR to use an interface VPC endpoint with an endpoint policy that includes the required permissions for images that the company needs to pull. Include a condition to allow the ECR operations for all account IDs in the company's organization.Schedule a daily Amazon EventBridge rule to invoke an AWS Lambda function that deletes all untagged images over the count of five.
**Answer: A (LEAVE A REPLY)**
Explanation: This option allows the company to use a private repository in Amazon ECR to store and manage its Docker images securely and efficiently1. By creating a permissions policy for the repository

that allows only required ECR operations, such as ecr:GetDownloadUrlForLayer, ecr:BatchGetImage, ecr:
BatchCheckLayerAvailability, ecr:PutImage, and ecr:InitiateLayerUpload2, the company can restrict access to the repository and prevent unauthorized actions. By including a condition to allow the ECR operations if the value of the aws:PrincipalOrgID condition key is equal to the ID of the company's organization, the company can ensure that only accounts that are within its organization can access the images3. By adding a lifecycle rule to the ECR repository that deletes all untagged images over the count of five, the company can reduce storage costs and retain only the most recent untagged images4.
:
Amazon ECR private repositories
Amazon ECR repository policies
Restricting access to AWS Organizations members
Amazon ECR lifecycle policies

**NEW QUESTION: 86**
A company is planning to migrate its on-premises transaction-processing application to AWS. The application runs inside Docker containers that are hosted on VMS in the company's data center. The Docker containers have shared storage where the application records transaction data.
The transactions are time sensitive. The volume of transactions inside the application is unpredictable. The company must implement a low-latency storage solution that will automatically scale throughput to meet increased demand. The company cannot develop the application further and cannot continue to administerthe Docker hosting environment.
How should the company migrate the application to AWS to meet these requirements?
**A.** Migrate the containers that run the application to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon S3 to store the transaction data that the containers share.
**B.** Migrate the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS). Create an Amazon Elastic File System (Amazon EFS) file system. Create a Fargate task definition. Add a volume to the task definition to point to the EFS file system
**C.** Migrate the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS). Create an Amazon Elastic Block Store (Amazon EBS) volume. Create a Fargate task definition. Attach the EBS volume to each running task.
**D.** Launch Amazon EC2 instances. Install Docker on the EC2 instances. Migrate the containers to the EC2 instances. Create an Amazon Elastic File System (Amazon EFS) file system. Add a mount point to the EC2 instances for the EFS file system.
**Answer: B (LEAVE A REPLY)**
Migrating the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS) will meet the requirement of not administering the Docker hosting environment. AWS Fargate is a serverless compute engine that runs containers without requiring any infrastructure management3.
Creating an Amazon Elastic File System (Amazon EFS) file system and adding a volume to the Fargate task definition to point to the EFS file system will meet the requirement of low-latency storage that will

automatically scale throughput to meet increased demand. Amazon EFS is a fully managed file system service that provides shared access to data from multiple containers, supports NFSv4 protocol, and offers consistent performance and high availability4. Amazon EFS also supports automatic scaling of throughput based on the amount of data stored in the file system5.

**NEW QUESTION: 87**
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 eventmessage 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.
**Answer: C,E,F (LEAVE A REPLY)**
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/lambda-edge-how-it-works-tutorial.
html

**NEW QUESTION: 88**
A solutions architect is reviewing a company's process for taking snapshots of Amazon RDS DB instances.
The company takes automatic snapshots every day and retains the snapshots for 7 days.
The solutions architect needs to recommend a solution that takes snapshots every 6 hours and retains the snapshots for 30 days. The company uses AWS Organizations to manage all of its AWS accounts.
The company needs a consolidated view of the health of the RDS snapshots.
Which solution will meet these requirements with the LEAST operational overhead?
**A.** Turn on the cross-account management feature in AWS Backup. Create a backup plan that specifies the frequency and retention requirements. Add a tag to the DB instances. Apply the backupplan by using tags. Use AWS Backup to monitor the status of the backups.
**B.** Turn on the cross-account management feature in Amazon RDS. Create a snapshot global policy that specifies the frequency and retention requirements. Use the RDS console in the management account to monitor the status of the backups.

**C.** Turn on the cross-account management feature in AWS CloudFormation. From the management account, deploy a CloudFormation stack set that contains a backup plan from AWS Backup that specifies the frequency and retention requirements. Create an AWS Lambda function in the management account to monitor the status of the backups. Create an Amazon EventBridge rule in each account to run the Lambda function on a schedule.
**D.** Configure AWS Backup in each account. Create an Amazon Data Lifecycle Manager lifecycle policy that specifies the frequency and retention requirements. Specify the DB instances as the target resource.Use the Amazon Data Lifecycle Manager console in each member account to monitor the status of the backups.
**Answer: (SHOW ANSWER)**
Turning on the cross-account management feature in AWS Backup will enable managing and monitoring backups across multiple AWS accounts that belong to the same organization in AWS Organizations1. Creating a backup plan that specifies the frequency and retention requirements will enable taking snapshots every 6 hours and retaining them for 30 days2. Adding a tag to the DB instances will enable applying the backup plan by using tags2. Using AWS Backup to monitor the status of the backups will enable having a consolidated view of the health of the RDS snapshots1.

**NEW QUESTION: 89**
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 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.
**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.
**Answer: D (LEAVE A REPLY)**

https://aws.amazon.com/premiumsupport/knowledge-center/vpc-analyze-inbound-traffic-nat-gateway/ by Cloudxie says "select appropriate log"

**NEW QUESTION: 90**
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.
**Answer: (SHOW ANSWER)**
This solution will meet the requirements with the least operational overhead because it allows the company to modify the existing environment's capacity configuration, so it becomes a load-balanced environment type.
By selecting all availability zones, the company can ensure that the application is running in multiple availability zones, which can help to improve the availability and scalability of the application. The company can also add a scale-out rule that will run if the average CPU utilization is over 85% for 5 minutes, which can help to mitigate the performance issues. This solution does not require creating new Elastic Beanstalk environments or rebuilding the existing one, which reduces the operational overhead.
You can refer to the AWS Elastic Beanstalk documentation for more information on how to use this service:
https://aws.amazon.com/elasticbeanstalk/You can refer to the AWS documentation for more information on how to use autoscaling:https://aws.amazon.com/autoscaling/

**NEW QUESTION: 91**
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, 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.
**B.** 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.
**C.** 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.
**D.** 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.
**Answer: B (LEAVE A REPLY)**



**NEW QUESTION: 92**
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
**Answer: A (LEAVE A REPLY)**
it will use the ALB to handle the unpredictable bursts of traffic and route it to the SQS queue. The SQS queue will act as a buffer to store incoming data temporarily and the container running in Amazon ECS with the Fargate launch type will process messages in the queue. This approach will ensure that all data is processed and prevent data loss.

**NEW QUESTION: 93**
A company hosts a blog post application on AWS using Amazon API Gateway, Amazon DynamoDB, and AWS Lambda. The application currently does not use API keys to authorize requests. The API model is as follows:
GET/posts/[postid] to get post details
GET/users[userid] to get user details
GET/comments/[commentid] to get comments details
The company has noticed users are actively discussing topics in the comments section, and the company wants to increase user engagement by marking the comments appears in real time.
Which design should be used to reduce comment latency and improve user experience?
**A.** Use edge-optimized API with Amazon CloudFront to cache API responses.
**B.** Modify the blog application code to request GET comment[commented] every 10 seconds.
**C.** Use AWS AppSync and leverage WebSockets to deliver comments.
**D.** Change the concurrency limit of the Lambda functions to lower the API response time.
**Answer: C (LEAVE A REPLY)**
https://docs.aws.amazon.com/appsync/latest/devguide/graphql-overview.html AWS AppSync is a fully managed GraphQL service that allows applications to securely access, manipulate, and receive data as well as real-time updates from multiple data sources1. AWS AppSync supports GraphQL subscriptions to perform real-time operations and can push data to clients that choose to listen to specific events from the backend1. AWS AppSync uses WebSockets to establish and maintain a secure connection between the clients and the API endpoint2. Therefore, using AWS AppSync and leveraging WebSockets is a suitable design to reduce comment latency and improve user experience.

**NEW QUESTION: 94**
A company has millions of objects in an Amazon S3 bucket. The objects are in the S3 Standard storage class.

All the S3 objects are accessed frequently. The number of users and applications that access the objects is increasing rapidly. The objects are encrypted with server-side encryption with AWS KMS Keys (SSE-KMS).
A solutions architect reviews the company's monthly AWS invoice and notices that AWS KMS costs are increasing because of the high number of requests from Amazon S3. The solutions architect needs to optimize costs with minimal changes to the application.
Which solution will meet these requirements with the LEAST operational overhead?
**A.** Create a new S3 bucket that has server-side encryption with customer-provided keys (SSE-C) as the encryption type. Copy the existing objects to the new S3 bucket. Specify SSE-C.
**B.** Create a new S3 bucket that has server-side encryption with Amazon S3 managed keys (SSE-S3) as the encryption type. Use S3 Batch Operations to copy the existing objects to the new S3 bucket. Specify SSE-S3.
**C.** Use AWS CloudHSM to store the encryption keys. Create a new S3 bucket. Use S3 Batch Operations to copy the existing objects to the new S3 bucket. Encrypt the objects by using the keys from CloudHSM.
**D.** Use the S3 Intelligent-Tiering storage class for the S3 bucket. Create an S3 Intelligent-Tiering archive configuration to transition objects that are not accessed for 90 days to S3 Glacier Deep Archive.
**Answer: B (LEAVE A REPLY)**
To reduce the volume of Amazon S3 calls to AWS KMS, use Amazon S3 bucket keys, which are protected encryption keys that are reused for a limited time in Amazon S3. Bucket keys can reduce costs for AWS KMS requests by up to 99%. You can configure a bucket key for all objects in an Amazon S3 bucket, or for a specific object in an Amazon S3
bucket.https://docs.aws.amazon.com/fr_fr/kms/latest/developerguide
/services-s3.html

**NEW QUESTION: 95**
A company needs to improve the reliability ticketing application. The application runs on an Amazon Elastic Container Service (Amazon ECS) cluster. The company uses Amazon CloudFront to servo the application. A single ECS service of the ECS cluster is the CloudFront distribution's origin.
The application allows only a specific number of active users to enter a ticket purchasing flow. These users are identified by an encrypted attribute in their JSON Web Token (JWT). All other users are redirected to a waiting room module until there is available capacity for purchasing.
The application is experiencing high loads. The waiting room modulo is working as designed, but load on the waiting room is disrupting the application's availability. This disruption is negatively affecting the application's ticket sale Transactions.
Which solution will provide the MOST reliability for ticket sale transactions during periods of high load? '
**A.** Create a separate service in the ECS cluster for the waiting room. Use a separate scaling configuration.
Ensure that the ticketing service uses the JWT info-nation and appropriately forwards requests to the waring room service.
**B.** Move the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Split the wailing room module into a pod that is separate from the ticketing pod. Make the ticketing pod part of a

StatefulSeL Ensure that the ticketing pod uses the JWT information and appropriately forwards requests to the waiting room pod.

**C.** Create a separate service in the ECS cluster for the waiting room. Use a separate scaling configuration.

Create a CloudFront function That inspects the JWT information and appropriatelyforwards requests to the ticketing service or the waiting room service

**D.** Move the application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Split the wailing room module into a pod that is separate from the ticketing pod. Use AWS App Mesh by provisioning the App Mesh controller for Kubermetes. Enable mTLS authentication and service-to-service authentication for communication between the ticketing pod and the waiting room pod. Ensure that the ticketing pod uses The JWT information and appropriately forwards requests to the waiting room pod.

**Answer: C (LEAVE A REPLY)**

Implementing a CloudFront function that inspects the JWT information and appropriately forwards requests either to the ticketing service or the waiting room service within the Amazon ECS cluster enhances reliability during high load periods. This solution segregates the load between the main application and the waiting room, ensuring that the ticketing service remains unaffected by the high load on the waiting room. Using CloudFront functions for request routing based on JWT attributes allows for efficient distribution of user traffic, thereby maintaining the application's availability and performance during peak times.

AWS Documentation on Amazon CloudFront Functions provides guidance on creating and deploying functions that can inspect and manipulate HTTP(S) requests at the edge, close to the users. This approach is in line with best practices for scaling and managing high-traffic web applications.

**NEW QUESTION: 96**

A company ingests and processes streaming market data. The data rate is constant. A nightly process that calculates aggregate statistics is run, and each execution takes about 4 hours to complete. The statistical analysis is not mission critical to the business, and previous data points are picked up on the next execution if a particular run fails.

The current architecture uses a pool of Amazon EC2 Reserved Instances with 1-year reservations running full time to ingest and store the streaming data in attached Amazon EBS volumes. On-Demand EC2 instances are launched each night to perform the nightly processing, accessing the stored data from NFS shares on the ingestion servers, and terminating the nightly processing servers when complete. The Reserved Instance reservations are expiring, and the company needs to determine whether to purchase new reservations or implement a new design.

Which is the most cost-effective design?

**A.** Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3. Use a scheduled script to launch a fleet of EC2 On-Demand Instances each night to perform the batch processing of the S3 data. Configure the script to terminate the instances when the processing is complete.

**B.** Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3. Use AWS Batch with Spot Instances to perform nightly processing with a maximum Spot price that is 50% of the On-Demand price.
**C.** Update the ingestion process to use a fleet of EC2 Reserved Instances with 3-year reservations behind a Network Load Balancer. Use AWS Batch with Spot Instances to perform nightly processing with a maximum Spot price that is 50% of the On-Demand price.
**D.** Update the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon Redshift.Use Amazon EventBridge to schedule an AWS Lambdafunction to run nightly to query Amazon Redshift to generate the daily statistics.
**Answer: (SHOW ANSWER)**
Updating the ingestion process to use Amazon Kinesis Data Firehose to save data to Amazon S3 will reduce the need for EC2 instances and EBS volumes for data storage1. Using AWS Batch with Spot Instances to perform nightly processing will leverage the cost savings of Spot Instances, which are up to 90% cheaper than On-Demand Instances2. AWS Batch will also handle the scheduling and scaling of the processing jobs.
Setting the maximum Spot price to 50% of the On-Demand price will reduce the chances of interruption and ensure that the processing is cost-effective.

**NEW QUESTION: 97**
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.
**Answer: B (LEAVE A REPLY)**
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.

**NEW QUESTION: 98**

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.
**Answer: B,C,F (LEAVE A REPLY)**
* Option A is incorrect because creating an SCP to set a fixed monthly account usage limit is not possible. SCPs are policies that specify the services and actions that users and roles can use in the member accounts of an AWS Organization. SCPscannot enforce budget limits or prevent users from launching costly services or running services unnecessarily1
* Option B is correct because using AWS Budgets to create a fixed monthly budget for each developer's account as part of the account creation process meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets allows you to plan your service usage, service costs, and instance reservations. You can create budgets that alert you when your costs or usage exceed (or are forecasted to exceed) your budgeted amount2
* Option C is correct because creating an SCP to deny access to costly services and components meets the requirement of ensuring that developers are not launching costly services or running services unnecessarily. SCPs can restrict access to certain AWS services or actions based on conditions such as region, resource tags, or request time. For example, an SCP can deny access to Amazon Redshift clusters or Amazon EC2 instances with certain instance types1
* Option D is incorrect because creating an IAM policy to deny access to costly services and components is not sufficient to meet the requirement of ensuring that developers are not launching costly services or running services unnecessarily. IAM policies can only control access to resources within a single AWS account. If developers have multiple accounts or can create new accounts, they can bypass the IAMpolicy restrictions. SCPs can apply across multiple accounts within an AWS Organization and prevent users from creating new accounts that do not comply with the SCP rules3

* Option E is incorrect because creating an AWS Budgets alert action to terminate services when the budgeted amount is reached is not possible. AWS Budgets alert actions can only perform one of the following actions: apply an IAM policy, apply an SCP, or send a notification through Amazon SNS. AWS Budgets alert actions cannot terminate services directly.
* Option F is correct because creating an AWS Budgets alert action to send an Amazon SNS notification when the budgeted amount is reached and invoking an AWS Lambda function to terminate all services meets the requirement of giving developers a fixed monthly budget to limit their AWS costs. AWS Budgets alert actions can send notifications through Amazon SNS when a budget threshold is breached. Amazon SNS can trigger an AWS Lambda function that can perform custom logic such as terminating all services in the developer's account. This way, developers cannot exceed their budget limit and incur additional costs.
References: 1:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html
2: https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/budgets-create.html
3: https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html
4: https://docs.aws.amazon.com/cost-management/latest/userguide/budgets-actions.html
5: https://docs.aws.amazon.com/sns/latest/dg/sns-lambda.html
6: https://docs.aws.amazon.com/lambda/latest/dg/welcome.html

**NEW QUESTION: 99**
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.** 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 astatic website Set up origin failover for the CloudFront distribution. Update the S3 static website to incorporate the custom error page.
**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.** 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.
**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.
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 100**
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? (Select 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.
**Answer: A,C (LEAVE A REPLY)**
https://aws.amazon.com/global-accelerator/faqs/

**NEW QUESTION: 101**
A company needs to implement disaster recovery for a critical application that runs in a single AWS Region.
The application's users interact with a web frontend that is hosted on Amazon EC2 Instances behind an Application Load Balancer (ALB). The application writes to an Amazon RD5 tor MySQL DB instance. The application also outputs processed documents that are stored in an Amazon S3 bucket The company's finance team directly queries the database to run reports. During busy periods, these queries consume resources and negatively affect application performance.
A solutions architect must design a solution that will provide resiliency during a disaster. The solution must minimize data loss and must resolve the performance problems that result from the finance team's queries.
Which solution will meet these requirements?
**A.** Migrate the database to Amazon DynamoDB and use DynamoDB global tables. Instruct the finance team to query a global table in a separate Region. Create an AWS Lambda function to periodically synchronize the contents of the original S3 bucket to a new S3 bucket in the separate Region. Launch EC2 instances and create an ALB in the separate Region. Configure the application to point to the new S3 bucket.
**B.** Launch additional EC2 instances that host the application in a separate Region. Add the additional instances to the existing ALB. In the separate Region, create a read replica of the RDS DB instance.

Instruct the finance team to run queries ageist the read replica. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 Docket in the separate Region. During a disaster, promote the read replace to a standalone DB instance. Configure the application to point to the new S3 bucket and to the newly project read replica.

**C.** Create a read replica of the RDS DB instance in a separate Region. Instruct the finance team to run queries against the read replica. Create AMIs of the EC2 instances mat host the application frontend- Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, promote the read replica to a standalone DB instance. Launch EC2 instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket.

**D.** Create hourly snapshots of the RDS DB instance. Copy the snapshots to a separate Region. Add an Amazon Elastic ache cluster m front of the existing RDS database. Create AMIs of the EC2 instances that host the application frontend Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, restore The database from the latest RDS snapshot. Launch EC2 Instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket

**Answer: C (LEAVE A REPLY)**

Implementing a disaster recovery strategy that minimizes data loss and addresses performance issues involves creating a read replica of the RDS DB instance in a separate region and directing the finance team's queries to this replica. This solution alleviates the performance impact on the primary database. Using Amazon S3 Cross- Region Replication (CRR) ensures that processed documents are available in the disaster recovery region. In the event of a disaster, the read replica can be promoted to a standalone DB instance, and EC2 instances can be launched from pre-created AMIs to serve the web frontend, thereby ensuring resiliency and minimal data loss.

AWS Documentation on Amazon RDS Read Replicas, Amazon S3 Cross-Region Replication, and Amazon EC2 AMIs provides comprehensive guidance on implementing a robust disaster recovery solution. This approach is in line with AWS best practices for high availability and disaster recovery planning.

**NEW QUESTION: 102**

A company wants to establish a dedicated connection between its on-premises infrastructure and AWS. The company is setting up a 1 Gbps AWS Direct Connect connection to its account VPC. The architecture includes a transit gateway and a Direct Connect gateway to connect multiple VPCs and the on-premises infrastructure.

The company must connect to VPC resources over a transit VIF by using the Direct Connect connection.

Which combination of steps will meet these requirements? (Select TWO.)

**A.** Update the 1 Gbps Direct Connect connection to 10 Gbps.

**B.** Advertise the on-premises network prefixes over the transit VIF.

**C.** Adverse the VPC prefixes from the Direct Connect gateway to the on-premises network over the transit VIF.

**D.** Update the Direct Connect connection's MACsec encryption mode attribute to must encrypt.
**E.** Associate a MACsec Connection Key Name-Connectivity Association Key (CKN/CAK) pair with the Direct Connect connection.
**Answer: B,C (LEAVE A REPLY)**
To connect VPC resources over a transit Virtual Interface (VIF) using a Direct Connect connection, the company should advertise the on-premises network prefixes over the transit VIF and advertise the VPC prefixes from the Direct Connect gateway to the on-premises network over the same VIF. This configuration ensures seamless connectivity between the on-premises infrastructure and the AWS VPCs through the transit gateway, facilitating efficient and secure communication across the network.
AWS Documentation on AWS Direct Connect and transit gateways provides detailed instructions on configuring transit VIFs and routing for Direct Connect connections. This setup is recommended in AWS best practices for establishing dedicated network connections between on-premises environments and AWS to achieve low-latency, high-throughput, and secure connectivity.

**NEW QUESTION: 103**
Question:
A company is running a large containerized workload in the AWS Cloud using Amazon ECS. The development team recently started usingAWS Fargateinstead of EC2 in the ECS cluster. The company is worried about reaching themaximum number of ECS tasksallowed in the account.
A solutions architect must implement a solution that notifies the development team when Fargate usage reaches80% of the quota.
What should the architect do?
**A.** Use CloudWatch to monitor the Sample Count for each service. Alert when usage exceeds 80%.
**B.** Use CloudWatch to monitor ECS service quotas under the AWS/Usage namespace. Create an alarm when utilization exceeds 80%. Notify via SNS.
**C.** Use a Lambda function to poll Fargate metrics. Notify via SES when usage exceeds 80%.
**D.** Use AWS Config to monitor Fargate quotas. Notify via SES if non-compliant.
**Answer: (SHOW ANSWER)**
Comprehensive and Detailed Explanation:
B is correct becauseAWS/Usagenamespace in CloudWatch provides built-in metrics for service quota usage.
You can create an alarm forFargate task count usageand notify the dev team usingSNSwhen it reaches 80% of the quota.
* A misuses "Sample Count" and doesn't relate to service quota.
* C is overly complex and not needed - native metrics exist.
* D misuses Config (meant for compliance, not utilization).
#Reference:
https://docs.aws.amazon.com/servicequotas/latest/userguide/monitoring-using-cloudwatch.html

**NEW QUESTION: 104**

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.
**Answer: B,D (LEAVE A REPLY)**
AWS Trusted Advisor is a service that provides real-time guidance to help users provision their resources following AWS best practices1. One of the Trusted Advisor checks is "Low Utilization Amazon EC2 Instances", which identifies EC2 instances that appear to be underutilized based on CPU, network I/O, and disk I/O metrics1. This check can help users optimize the cost and size of their EC2 instances by recommending smaller or more appropriate instance types.
AWS Compute Optimizer is a service that analyzes the configuration and utilization metrics of AWS resources and generates optimization recommendations to reduce the costand improve the performance of workloads2. Compute Optimizer supports four types of AWS resources: EC2 instances, EBS volumes, ECS services on AWSFargate, and Lambda functions2. For EC2 instances, Compute Optimizer evaluates the vCPUs, memory, storage, and other specifications, as well as the CPU utilization, network in and out, disk read and write, and other utilization metrics of currently running instances3. It then recommends optimal instance types based on price-performance trade-offs.
Option A is incorrect because purchasing AWS Business Support or AWS Enterprise Support for the account will not directly help with cost-optimization and sizing of EC2 instances. However, these support plans do provide access to more Trusted Advisor checks than the basic support plan1.
Option C is incorrect because installing the Amazon CloudWatch agent and configuring memory metric collection on the EC2 instances will not provide any optimization recommendations by itself. However, memory metrics can be used by Compute Optimizer to enhance its recommendations if enabled3.
Option E is incorrect because creating an EC2 Instance Savings Plan for the AWS Regions, instance families, and operating systems of interest will not help with cost-optimization and sizing of EC2 instances. Savings Plans are a flexible pricing model that offer lower prices on Amazon EC2 usage in exchange for a commitment to a consistent amount of usage for a 1- or 3-year term4. Savings Plans do not affect the configuration or utilization of EC2 instances.

**NEW QUESTION: 105**
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 on-going administration.
* A visualization toot to create dashboards to observe events in near-real time.
* Support for semi -structured JSON data and dynamic schemas.
Which combination of components will enabled company to create a monitoring solution that will satisfy these requirements" (Select TWO.)
**A.** Use Amazon Kinesis Data Firehose to buffer events Create an AWS Lambda function 10 process and transform events
**B.** Create an Amazon Kinesis data stream to buffer events Create an AWS Lambda function to process and transform evens
**C.** Configure an Amazon Aurora PostgreSQL DB cluster to receive events Use Amazon Quick Sight 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 0 DB instance to receive events Use Amazon QuickSight to read from the database and create near-real-time visualizations and dashboards
**Answer: A,D (LEAVE A REPLY)**
https://aws.amazon.com/kinesis/data-firehose/faqs/

**NEW QUESTION: 106**
Question:
A company provisions short-lived AWS accounts for students. Each account needs access to ml.p2.xlarge SageMaker instances for training and inference. The default quotas are insufficient.
How should quota increases be automated during account provisioning?
**A.** Create a quota request template inus-east-1, enable template association, and add quotas for ml.p2. xlarge training and endpoint usage in ap-southeast-2.
**B.** Use ml.p2.xlarge training warm pool quota in ap-southeast-2.
**C.** Create the template in ap-southeast-2 for SageMaker quotas in us-east-1.
**D.** Use warm pool quotas in us-east-1.
**Answer: A (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
A is correct:Quota request templates must be created in us-east-1, the only region that supports them. You specify the desired quotain the correct target Region(ap-southeast-2) for SageMaker training and endpoint usage. This enablesautomatic quota increasesfor newly created accounts in the org.
* B, C, and D misconfigure either region or quota type.
#Reference:Quota request templates

**NEW QUESTION: 107**
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 to 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.
**Answer: (SHOW ANSWER)**
This solution allows for deploying the Lambda function and API Gateway endpoint to another region, providing a failover option in case of any issues in the primary region. Using Route 53's failover routing policy allows for automatic routing of traffic to the healthy endpoint, ensuring that the application is available even in case of issues in one region. This solution provides a cost-effective and simple way to implement failover while minimizing operational overhead.

**NEW QUESTION: 108**
A company needs to use an AWS Transfer Family SFTP-enabled server with an Amazon S3 bucket to receive updates from a third-party data supplier. The data is encrypted with Pretty Good Privacy (PGP) encryption The company needs a solution that will automatically decrypt the data after the company receives the data A solutions architect will use a Transfer Family managed workflow The company has created an 1AM service role by using an 1AM policy that allows access to AWS Secrets Manager and the S3 bucket The role's trust relationship allows the transfer amazonaws com service to assume the rote What should the solutions architect do next to complete the solution for automatic decryption'?

**A.** Store the PGP public key in Secrets Manager Add a nominal step in the Transfer Family managed workflow to decrypt files Configure PGP encryption parameters in the nominal step Associate the workflow with the Transfer Family server
**B.** Store the PGP private key in Secrets Manager Add an exception-handling step in the Transfer Family managed workflow to decrypt files Configure PGP encryption parameters in the exception handler Associate the workflow with the SFTP user
**C.** Store the PGP private key in Secrets Manager Add a nominal step in the Transfer Family managed workflow to decrypt files. Configure PGP decryption parameters in the nominal step Associate the workflow with the Transfer Family server
**D.** Store the PGP public key in Secrets Manager Add an exception-handling step in the Transfer Family managed workflow to decrypt files Configure PGP decryption parameters in the exception handler Associate the workflow with the SFTP user
**Answer: (SHOW ANSWER)**
* Store the PGP Private Key:
* Step 1: In the AWS Management Console, navigate to AWS Secrets Manager.
* Step 2: Store the PGP private key in Secrets Manager. Ensure the key is encrypted and properly secured.
* Set Up the Transfer Family Managed Workflow:
* Step 1: In the AWS Transfer Family console, create a new managed workflow.
* Step 2: Add a nominal step to the workflow that includes the decryption of the files. Configure this step with the PGP decryption parameters, referencing the PGP private key stored in Secrets Manager.
* Step 3: Associate this workflow with the Transfer Family SFTP server, ensuring that incoming files are automatically decrypted upon receipt.
This solution ensures that the data is securely decrypted as it is transferred from the SFTP server to the S3 bucket, automating the decryption process and leveraging AWS Secrets Manager for key management.
References
* AWS Transfer Family Documentation
* Using AWS Secrets Manager for Managing Secrets
* AWS Transfer Family Managed Workflows

**NEW QUESTION: 109**
A company that provisions job boards for a seasonal workforce is seeing an increase in traffic and usage. The backend services run on a pair of Amazon EC2 instances behind an Application Load Balancer with Amazon DynamoDB as the datastore. Application read and write traffic is slow during peak seasons.
Which option provides a scalable application architecture to handle peak seasons with the LEAST development effort?
**A.** Migrate the backend services to AWS Lambda. Increase the read and write capacity of DynamoDB.
**B.** Migrate the backend services to AWS Lambda. Configure DynamoDB to use global tables.
**C.** Use Auto Scaling groups for the backend services. Use DynamoDB auto scaling.

**D.** Use Auto Scaling groups for the backend services. Use Amazon Simple Queue Service (Amazon SQS) and an AWS Lambda function to write to DynamoDB.

**Answer: C (LEAVE A REPLY)**

* Option C is correct because using Auto Scaling groups for the backend services allows the company to scale up or down the number of EC2 instances based on the demand and traffic. This way, the backend services can handle more requests during peak seasons without compromising performance or availability. Using DynamoDB auto scaling allows the company to adjust the provisioned read andwrite capacity of the table or index automatically based on the actual traffic patterns. This way, the table or index can handle sudden increases or decreases in workload without throttling or overprovisioning1.
* Option A is incorrect because migrating the backend services to AWS Lambda may require significant development effort to rewrite the code and test the functionality. Moreover, increasing the read and write capacity of DynamoDB manually may not be efficient or cost-effective, as it does not account for the variability of the workload. The company may end up paying for unused capacity or experiencing throttling if the workload exceeds the provisioned capacity1.
* Option B is incorrect because migrating the backend services to AWS Lambda may require significant development effort to rewrite the code and test the functionality. Moreover, configuring DynamoDB to use global tables may not be necessary or beneficial for the company, as global tables are mainly used for replicating data across multiple AWS Regions for fast local access and disaster recovery. Global tables do not automatically scale the provisioned capacity of each replica table; they still require manual or auto scaling settings2.
* Option D is incorrect because using Amazon Simple Queue Service (Amazon SQS) and an AWS Lambda function to write to DynamoDB may introduce additional complexity and latency to the application architecture. Amazon SQS is a message queue service that decouples and coordinates the components of a distributed system. AWS Lambda is a serverless compute service that runs code in response to events. Using these services may require significant development effort to integrate them with the backend services and DynamoDB. Moreover, they may not improve the read performance of DynamoDB, which may also be affected by high traffic3.

:

Auto Scaling groups
DynamoDB auto scaling
AWS Lambda
DynamoDB global tables
AWS Lambda vs EC2: Comparison of AWS Compute Resources - Simform
Managing throughput capacity automatically with DynamoDB auto scaling - Amazon DynamoDB AWS Aurora Global Database vs. DynamoDB Global Tables Amazon Simple Queue Service (SQS)

**NEW QUESTION: 110**

A company has an loT platform that runs in an on-premises environment. The platform consists of a server that connects to loT 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? (Select 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
**Answer: B,D,E (LEAVE A REPLY)**
https://aws.amazon.com/step-functions/use-cases/

**NEW QUESTION: 111**
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.
**Answer: D (LEAVE A REPLY)**
By breaking down the monolithic API into individual Lambda functions and using API Gateway to handle the incoming requests, the solution can automatically scale to handle the new and varying load without the need for manual scaling actions. Additionally, this option will automatically handle the traffic without

the need of having EC2 instances running all the time and only pay for the number of requests and the duration of the execution of the Lambda function.
By updating the Route 53 record to point to the API Gateway, the solution can handle the traffic and also it will direct the traffic to the correct endpoint.

**NEW QUESTION: 112**
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 OuickSight integration with OpsWorks to generate patch compliance reports.
**C.** Use an Amazon EventBridge (Amazon CloudWatch Events) 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.
**Answer: A (LEAVE A REPLY)**
https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-patch.html

**NEW QUESTION: 113**
A company is planning to migrate an application from on premises to the AWS Cloud The company will begin the migration by moving the application underlying data storage to AWS The application data is stored on a shared tile system on premises and the application servers connect to the shared file system through SMB A solutions architect must implement a solution that uses an Amazon S3 bucket for shared storage. Until the application is fully migrated and code is rewritten to use native Amazon S3 APIs the application must continue to have access to the data through SMB The solutions architect must migrate the application data to AWS (o its new location while still allowing the on-premises application to access the data Which solution will meet these requirements?
**A.** Create a new Amazon FSx for Windows File Server file system Configure AWS DataSync with one location for the on-premises file share and one location for the new Amazon FSx file system Create a new DataSync task to copy the data from the on-premises file share location to the Amazon FSx file system
**B.** Create an S3 bucket for the application Copy the data from the on-premises storage to the S3 bucket
**C.** Deploy an AWS Server Migration Service (AWS SMS) VM to the on-premises environment Use AWS SMS to migrate the file storage server from on premises to an Amazon EC2 instance
**D.** Create an S3 bucket for the application Deploy a new AWS Storage Gateway file gateway on an on-premises VM Create a new file share that stores data in the S3 bucket and is associated with the file gateway Copy the data from the on-premises storage to the new file gateway endpoint

**Answer: D (LEAVE A REPLY)**

* Create an S3 Bucket:
* Log in to the AWS Management Console and navigate to Amazon S3.
* Create a new S3 bucket that will serve as the destination for the application data.
* Deploy AWS Storage Gateway:
* Download and deploy the AWS Storage Gateway virtual machine (VM) on your on-premises environment. This VM can be deployed on VMware ESXi, Microsoft Hyper-V, or Linux KVM.
* Configure the File Gateway:
* Configure the deployed Storage Gateway as a file gateway. This will enable it to present Amazon S3 buckets as SMB file shares to your on-premises applications.
* Create a New File Share:
* Within the Storage Gateway configuration, create a new file share that is associated with the S3 bucket you created earlier. This file share will use the SMB protocol, allowing your on-premises applications to access the S3 bucket as if it were a local SMB file share.
* Copy Data to the File Gateway:
* Use your preferred method (such as robocopy, rsync, or similar tools) to copy data from the on-premises storage to the newly created file gateway endpoint. This data will be stored in the S3 bucket, maintaining accessibility through SMB.
* Ensure Secure and Efficient Data Transfer:
* AWS Storage Gateway ensures that all data in transit is encrypted using TLS, providing secure data transfer to AWS. It also provides local caching for frequently accessed data, improving access performance for on-premises applications.

This approach allows your existing on-premises applications to continue accessing data via SMB while leveraging the scalability and durability of Amazon S3.

References

* AWS Storage Gateway Overview#67#.
* AWS DataSync and Storage Gateway Hybrid Architecture#66#.
* AWS S3 File Gateway Details#68#.

**NEW QUESTION: 114**

A company is using AWS CodePipeline for the CI/CD of an application to an Amazon EC2 Auto Scaling group. All AWS resources are defined in AWS CloudFormation templates. The application artifacts are stored in an Amazon S3 bucket and deployed to the Auto Scaling group using instance user data scripts. As the application has become more complex, recent resource changes in the CloudFormation templates have caused unplanned downtime.

How should a solutions architect improve the CI/CD pipeline to reduce the likelihood that changes in the templates will cause downtime?

**A.** Adapt the deployment scripts to detect and report CloudFormation error conditions when performing deployments. Write test plans for a testing team to execute in a non-production environment before approving the change for production.

**B.** Use plugins for the integrated development environment (IDE) to check the templates for errors, and use the AWS CLI to validate that the templates are correct. Adapt the deployment code to check for error conditions and generate notificationson errors. Deploy to a test environment and execute a manual test plan before approving the change for production.
**C.** Implement automated testing using AWS CodeBuild in a test environment. Use CloudFormation change sets to evaluate changes before deployment. Use AWS CodeDeploy to leverage blue/green deployment patterns to allow evaluations and the ability to revert changes, if needed.
**D.** Use AWS CodeDeploy and a blue/green deployment pattern with CloudFormation to replace the user data deployment scripts. Have the operators log in to running instances and go through a manual test plan to verify the application is running as expected.
**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 115**
A company is rearchitecting its applications to run on AWS. The company's infrastructure includes multiple Amazon EC2 instances. The company's development team needs different levels of access. The company wants to implement a policy that requires all Windows EC2 instances to be joined to an Active Directory domain on AWS. The company also wants to Implement enhanced security processes such as multi-factor authentication (MFA). The company wants to use managed AWS services wherever possible. Which solution will meet these requirements?
**A.** Create an AWS Directory Service for Microsoft Active Directory implementation. Launch an Amazon Workspace. Connect to and use the Workspace for domain security configuration tasks.
**B.** Create an AWS Directory Service for Microsoft Active Directory implementation. Launch an EC2 instance. Connect to and use the EC2 instance for domain security configuration tasks.
**C.** Create an AWS Directory Service Simple AD implementation. Launch an EC2 instance. Connect to and use the EC2 instance for domain security configuration tasks.
**D.** Create an AWS Directory Service Simple AD implementation. Launch an Amazon Workspace.Connect to and use the Workspace for domain security configuration tasks.
**Answer: A (LEAVE A REPLY)**
A is the correct answer because it uses AWS Directory Service for Microsoft Active Directory to join the Windows EC2 instances to an Active Directory domain on AWS and enable MFA. AWS Directory Service for Microsoft Active Directory, also known as AWS Managed Microsoft AD, is a fully managed service that is powered by Windows Server 2019. It allows you to run directory-aware workloads in the AWS Cloud, including Microsoft SharePoint and custom .NET and SQL Server-based applications. You can also configure a trust relationship between AWS Managed Microsoft AD in the AWS Cloud and your existing on-premises Microsoft Active Directory. AWS Managed Microsoft AD supports MFA by integrating with your existing RADIUS-based MFA infrastructure. To join the Windows EC2 instances to an Active Directory domain on AWS, you can use an Amazon Workspace, which is a fully managed, secure desktop computing service that runs on AWS. You can connect to and use the Workspace for domain security configuration tasks. References:
* https://docs.aws.amazon.com/directoryservice/latest/admin-guide/directory_microsoft_ad.html
* https://docs.aws.amazon.com/directoryservice/latest/admin-guide/ms_ad_join_instance.html

* https://docs.aws.amazon.com/workspaces/latest/adminguide/amazon-workspaces.html

**NEW QUESTION: 116**
A company operates quick-service restaurants. The restaurants follow a predictable model with high sales traffic for 4 hours daily Sales traffic is lower outside of those peak hours.
The point of sale and management platform is deployed in the AWS Cloud and has a backend that is based on Amazon DynamoDB. The database table uses provisioned throughput mode with 100.000 RCUs and 80.000 WCUs to match known peak resource consumption.
The company wants to reduce its DynamoDB cost and minimize the operational overhead for the IT staff.
Which solution meets these requirements MOST cost-effectively?
**A.** Reduce the provisioned RCUs and WCUs
**B.** Change the DynamoDB table to use on-demand capacity.
**C.** Enable Dynamo DB auto scaling tor the table
**D.** Purchase 1-year reserved capacity that is sufficient to cover the peak load for 4 hours each day.
**Answer: (SHOW ANSWER)**
https://aws.amazon.com/blogs/database/amazon-dynamodb-auto-scaling-performance-and-cost-optimization- at-any-scale/ "As you can see, there are compelling reasons to use DynamoDB auto scaling with actively changing traffic. Auto scaling responds quickly and simplifies capacity management, which lowers costs by scaling your table's provisionedcapacity and reducing operational overhead."

**NEW QUESTION: 117**
A live-events company is designing a scaling solution for its ticket application on AWS. The application has high peaks of utilization during sale events. Each sale event is a one-time event that is scheduled. The application runs on Amazon EC2 instances that are in an Auto Scaling group. The application uses PostgreSOL for the database layer.
The company needs a scaling solution to maximize availability during the sale events.
Which solution will meet these requirements?
**A.** Use a predictive scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSOL Serverless v2 Multi-AZ DB instance with automatically scaling read replicas. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale event. Create an Amazon EventBridge rule to invoke the state machine.
**B.** Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazcyl ROS for PostgreSQL Multi-AZ DB instance with automatically scaling read replicas. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger read replica before a sale event. Fail over to the larger read replica. Create another EventBridge rule that invokes another Lambda function to scale down the read replica after the sale event.
**C.** Use a predictive scaling policy for the EC2 instances. Host the database on an Amazon RDS for PostgreSOL Multi-AZ DB instance with automatically scaling read replica. Create an AWS Step Functions state machine to run parallel AWS Lambda functions to pre-warm the database before a sale event. Create an Amazon EventBridge rule to invoke the state machine.

**D.** Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSQL Multi-AZ DB duster. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale event. Fail over to the larger Aurora Replica.Create another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event.

**Answer: D (LEAVE A REPLY)**

The correct answer is D. Use a scheduled scaling policy for the EC2 instances. Host the database on an Amazon Aurora PostgreSQL Multi-AZ DB cluster. Create an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale event. Fail over to the larger Aurora Replica.

Create another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event.

This solution will meet the requirements of maximizing availability during the sale events. A scheduled scaling policy for the EC2 instances will allow the application to scale up and down according to the predefined schedule of the sale events. Hosting the database on an Amazon Aurora PostgreSQL Multi-AZ DB cluster will provide high availability and durability, as well as compatibility with PostgreSQL. Creating an Amazon EventBridge rule that invokes an AWS Lambda function to create a larger Aurora Replica before a sale event will ensure that the database can handle theincreased read traffic during the peak periods. Failing over to the larger Aurora Replica will make it the primary instance, which will also improve the write performance of the database. Creating another EventBridge rule that invokes another Lambda function to scale down the Aurora Replica after the sale event will reduce the cost and resources of the database.

Reference: [3], section "Scaling Amazon Aurora MySQL and PostgreSQL with Aurora Auto Scaling"

**NEW QUESTION: 118**

An AWS partner company is building a service in AWS Organizations using Its organization named org. 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 privilegesecurity access using an API or command line tool to the customer account What is the MOST secure way to allow org1 to access resources h 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 rote's Amazon Resource Name (ARN) when requesting access to perform the required tasks

**D.** The customer should create an IAM rote and assign the required permissions to the IAM rote. The partner company should then use the IAM rote's Amazon Resource Name (ARN). Including the external ID in the IAM role's trust pokey, when requesting access to perform the required tasks

**Answer: (SHOW ANSWER)**

https://docs.aws.amazon.com/IAM/latest/UserGuide/confused-deputy.html

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.

**NEW QUESTION: 119**

A company runs its application on Amazon EC2 instances and AWS Lambda functions. The EC2 instances experience a continuous and stable load. The Lambda functions experience a varied and unpredictable load. The application includes a cachinglayer that uses an Amazon MemoryDB for Redis cluster.

A solutions architect must recommend a solution to minimize the company's overall monthly costs.

Which solution will meet these requirements?

**A.** Purchase an EC2 Instance Savings Plan to cover the EC2 instances. Purchase a Compute Savings Plan for Lambda to cover the minimum expectedconsumption of the Lambda functions. Purchase reserved nodes to cover the MemoryDB cache nodes.

**B.** Purchase a Compute Savings Plan to cover the EC2 instances. Purchase Lambda reserved concurrency to cover the expected Lambda usage. Purchasereserved nodes to cover the MemoryDB cache nodes.

**C.** Purchase a Compute Savings Plan to cover the entire expected cost of the EC2 instances, Lambda functions, and MemoryDB cache nodes.

**D.** Purchase a Compute Savings Plan to cover the EC2 instances and the MemoryDB cache nodes.Purchase Lambda reserved concurrency to cover theexpected Lambda usage.

**Answer: (SHOW ANSWER)**

This option uses different types of savings plans and reserved nodes to minimize the company's overall monthly costs for running its application on EC2 instances, Lambda functions, and MemoryDB cache nodes.

Savings plans are flexible pricing models that offer significant savings on AWS usage (up to 72%) in exchange for a commitment of a consistent amount of usage (measured in $/hour) for a one-year or three-year term. There are two types of savings plans: Compute Savings Plans and EC2 Instance Savings Plans.

Compute Savings Plans apply to any compute usage across EC2 instances, Fargate containers, Lambda functions, SageMaker notebooks, and ECS tasks. EC2 Instance Savings Plans apply to a specific instance family within a region and provide more savings than Compute Savings Plans (up to 66% versus up to 54%).

Reserved nodes are similar to savings plans but apply only to MemoryDB cache nodes. They offer up to 55% savings compared to on-demand pricing.

**NEW QUESTION: 120**
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 Account.
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? (Select 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
**Answer: A,C,F (LEAVE A REPLY)**
https://aws.amazon.com/premiumsupport/knowledge-center/cross-account-access-denied-error-s3/

**NEW QUESTION: 121**
A company that provides image storage services wants to deploy a customer-lacing solution to AWS. Millions of individual customers will use the solution. The solution will receive batches of large image files, resize the files, and store the files in an Amazon S3 bucket for up to 6 months.
The solution must handle significant variance in demand. The solution must also be reliable at enterprise scale and have the ability to rerun processing jobs in the event of failure.
Which solution will meet these requirements MOST cost-effectively?
**A.** Use AWS Step Functions to process the S3 event that occurs when a user stores an image. Run an AWS Lambda function that resizes the image in place and replaces the original file in the S3 bucket. Create an S3 Lifecycle expiration policy to expire all stored images after 6 months.
**B.** Use Amazon EventBridge to process the S3 event that occurs when a user uploads an image. Run an AWS Lambda function that resizes the image in place and replaces the original file in the S3 bucket.

Create an S3 Lifecycle expiration policy to expire all stored images after 6 months.

**C.** Use S3 Event Notifications to invoke an AWS Lambda function when a user stores an image. Use the Lambda function to resize the image in place and to store the original file in the S3 bucket. Create an S3 Lifecycle policy to move all stored images to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months.

**D.** Use Amazon Simple Queue Service (Amazon SQS) to process the S3 event that occurs when a user stores an image. Run an AWS Lambda function that resizes the image and stores the resized file in an S3 bucket that uses S3 Standard-Infrequent Access (S3 Standard-IA). Create an S3 Lifecycle policy to move all stored images to S3 Glacier Deep Archive after 6 months.

**Answer: C (LEAVE A REPLY)**

S3 Event Notifications is a feature that allows users to receive notifications when certain events happen in an S3 bucket, such as object creation or deletion1. Users can configure S3 Event Notifications to invoke an AWS Lambda function when a user stores an image in the bucket. Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources2. The Lambda function can resize the image in place and store the original file in the same S3 bucket. This way, the solution can handle significant variance in demand and be reliable at enterprise scale. Thesolution can also rerun processing jobs in the event of failure by using the retry and dead-letter queue features of Lambda2.

S3 Lifecycle is a feature that allows users to manage their objects so that they are stored cost-effectively throughout their lifecycle3. Users can create an S3 Lifecycle policy to move all stored images to S3 Standard- Infrequent Access (S3 Standard-IA) after 6 months. S3 Standard-IA is a storage class designed for data that is accessed less frequently, but requires rapid access when needed4. It offers a lower storage cost than S3 Standard, but charges a retrieval fee. Therefore, moving the images to S3 Standard-IA after 6 months can reduce the storage cost for the solution.

Option A is incorrect because using AWS Step Functions to process the S3 event that occurs when a user stores an image is not necessary or cost-effective. AWS Step Functions is a service that lets users coordinate multiple AWS services into serverless workflows. However, for this use case, a single Lambda function can handle the image resizing task without needing Step Functions.

Option B is incorrect because using Amazon EventBridge to process the S3 event that occurs when a user uploads an image is not necessary or cost-effective. Amazon EventBridge is a serverless event bus service that makes it easy to connect applications with data from a variety of sources. However, for this use case, S3 Event Notifications can directly invoke the Lambda function without needing EventBridge.

Option D is incorrect because using Amazon Simple Queue Service (Amazon SQS) to process the S3 event that occurs when a user stores an image is not necessary or cost-effective. Amazon SQS is a fully managed message queuing service that enables users to decouple and scale microservices, distributed systems, and serverless applications. However, for this use case, S3 Event Notifications can directly invoke the Lambda function without needing SQS. Moreover, storing the resized file in an S3 bucket that uses S3 Standard-IA will incur a retrieval fee every time the file is accessed, which may not be cost-effective for frequently accessed files.

**NEW QUESTION: 122**
A company is running an application on Amazon EC2 instances in the AWS Cloud. The application is using a MongoDB database with a replica set as its data tier. The MongoDB database is installed on systems in the company's on-premises data center and is accessible through an AWS Direct Connect connection to the data center environment.
A solutions architect must migrate the on-premises MongoDB database to Amazon DocumentDB (with MongoDB compatibility).
Which strategy should the solutions architect choose to perform this migration?
**A.** Create a fleet of EC2 instances. Install MongoDB Community Edition on the EC2 instances, and create a database. Configure continuous synchronous replication with the database that is running in the on- premises data center.
**B.** Create an AWS Database Migration Service (AWS DMS) replication instance. Create a source endpoint for the on-premises MongoDB database by using change data capture (CDC). Create a target endpoint for the Amazon DocumentDB database. Create and run a DMS migration task.
**C.** Create a data migration pipeline by using AWS Data Pipeline. Define data nodes for the on-premises MongoDB database and the Amazon DocumentDB database. Create a scheduled task to run the data pipeline.
**D.** Create a source endpoint for the on-premises MongoDB database by using AWS Glue crawlers.Configure continuous asynchronous replication between the MongoDB database and the Amazon DocumentDB database.
**Answer: B (LEAVE A REPLY)**
https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-mongodb-to-documentdb/

**NEW QUESTION: 123**
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 communicatewith 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.
**Answer: C (LEAVE A REPLY)**
Comprehensive and Detailed in Depth Explanation:
C is correct because AWS Transit Gateway is the most scalable and efficient way to interconnect hundreds of VPCs. By deploying one transit gateway per OU and sharing it with AWS RAM, each OU can isolate its network traffic and maintain internal communication without affecting or exposing other OUs.
References:
AWS Transit Gateway Best Practices
Using RAM with Transit Gateway

**NEW QUESTION: 124**
A company owns a chain of travel agencies and is running an application in the AWS Cloud. Company employees use the application to search for information about travel destinations. Destination content is updated four times each year.
Two fixed Amazon EC2 instances serve the application. The company uses an Amazon Route 53 public hosted zone with a multivalue record of travel.example.com that returns the Elastic IP addresses for the EC2 instances. The application uses Amazon DynamoDB as its primary data store. The company uses a self-hosted Redis instance as a caching solution.
During content updates, the load on the EC2 instances and the caching solution increases drastically. This increased load has led to downtime on several occasions. A solutions architect must update the application so that the application is highly available and can handle the load that is generated by the content updates.
Which solution will meet these requirements?
**A.** Set up DynamoDB Accelerator (DAX) as in-memory cache. Update the application to use DAX. Create an Auto Scaling group for the EC2 instances. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias. Configure scheduled scaling for the EC2 instances before the content updates.
**B.** Set up Amazon ElastiCache for Redis. Update the application to use ElastiCache. Create an Auto Scaling group for the EC2 instances. Create an Amazon CloudFront distribution, and set the Auto Scaling group as an origin for the distribution. Update the Route 53 record to use a simple routing policy that

targets the CloudFront distribution's DNS alias. Manually scale up EC2 instances before the content updates.
**C.** Set up Amazon ElastiCache for Memcached. Update the application to use ElastiCache Create an Auto Scaling group for the EC2 instances. Create an Application Load Balancer (ALB). Set the Auto Scaling group as a target for the ALB. Update the Route 53 record to use a simple routing policy that targets the ALB's DNS alias. Configure scheduled scaling for the application before the content updates.
**D.** Set up DynamoDB Accelerator (DAX) as in-memory cache. Update the application to use DAX. Create an Auto Scaling group for the EC2 instances. Create an Amazon CloudFront distribution,and set the Auto Scaling group as an origin for the distribution. Update the Route 53 record to use a simple routing policy that targets the CloudFront distribution's DNS alias. Manually scale up EC2 instances before the content updates.
**Answer: A (LEAVE A REPLY)**
Explanation: This option allows the company to use DAX to improve the performance and reduce the latency of the DynamoDB queries by caching the results in memory1. By updating the application to use DAX, the company can reduce the load on the DynamoDB tables and avoid throttling errors1. By creating an Auto Scaling group for the EC2 instances, the company can adjust the number of instances based on the demand and ensure high availability2. By creating an ALB, the company can distribute the incoming traffic across multiple EC2 instances and improve fault tolerance3. By updating the Route 53 record to use a simple routing policy that targets the ALB's DNS alias, the company can route users to the ALB endpoint and leverage its health checks and load balancing features4. By configuring scheduled scaling for the EC2 instances before the content updates, the company can anticipate and handle traffic spikes during peak periods5.
:
What is Amazon DynamoDB Accelerator (DAX)?
What is Amazon EC2 Auto Scaling?
What is an Application Load Balancer?
Choosing a routing policy
Scheduled scaling for Amazon EC2 Auto Scaling

**NEW QUESTION: 125**
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 mam 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 mam AWS account Establish connectivity between the Client VPN endpoint and the AWS Site-to-Site VPN
**Answer: (SHOW ANSWER)**
https://docs.aws.amazon.com/vpn/latest/clientvpn-admin/scenario-peered.html

**NEW QUESTION: 126**
A company hosts a web application on AWS in the us-east-1 Region The application servers are distributed across three Availability Zones behind an Application Load Balancer. The database is hosted in a MySQL database on an Amazon EC2 instance A solutions architect needs to design a Cross-Region data recovery solution using AWS services with an RTO of less than 5 minutes and an RPO of less than 1 minute. The solutions architect is deploying application servers in us-west-2, and has configured Amazon Route 53 hearth checks and DNS failover to us-west-2 Which additional step should the solutions architect take?
**A.** Migrate the database to an Amazon RDS tor MySQL instance with a cross-Region read replica in us-west-2
**B.** Migrate the database to an Amazon Aurora global database with the primary in us-east-1 and the secondary in us-west-2
**C.** Migrate the database to an Amazon RDS for MySQL instance with a Multi-AZ deployment.
**D.** Create a MySQL standby database on an Amazon EC2 instance in us-west-2
**Answer: (SHOW ANSWER)**
https://aws.amazon.com/rds/aurora/global-database/

**NEW QUESTION: 127**
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 ate 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.examWe.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 sflp.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.
**Answer: (SHOW ANSWER)**

https://aws.amazon.com/aws-transfer-family/faqs/
https://docs.aws.amazon.com/transfer/latest/userguide/what-is-aws-transfer-family.html
https://aws.amazon.com/about-aws/whats-new/2018/11/aws-transfer-for-sftp-fully-managed-sftp-for-s3/?nc1=h_ls

**NEW QUESTION: 128**
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 anAWS 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.
**Answer: B (LEAVE A REPLY)**
https://aws.amazon.com/blogs/security/how-to-securely-provide-database-credentials-to-lambda-functions-by- using-aws-secrets-manager/
https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html
https://docs.aws.amazon.com/secretsmanager/latest/userguide/integrating_cloudformation.html

**NEW QUESTION: 129**
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 ScaSng group The web application stores all blog content on an Amazon EFS volume.
The company recently added a feature 'or Moggers 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 win resolve the issues for users?

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

**B.** Update the Nog site to use instance store volumes tor storage. Copy the site contents to the volumes at launch and to Amazon S3 al 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.

**Answer: C (LEAVE A REPLY)**

https://aws.amazon.com/premiumsupport/knowledge-center/cloudfront-https-connection-fails/ Using an Amazon S3 bucket Using a MediaStore container or a MediaPackage channel Using an Application Load Balancer Using a Lambda function URL Using Amazon EC2 (or another custom origin) Using CloudFront origin groups

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

**NEW QUESTION: 130**

A company hosts a VPN in an on-premises data center. Employees currently connect to the VPN to access files in their Windows home directories. Recently, there has been a large growth in the number of employees who work remotely. As a result, bandwidth usage for connections into the data center has begun to reach

100% during business hours.

The company must design a solution on AWS that will support the growth of the company's remote workforce, reduce the bandwidth usage for connections into the data center, and reduce operational overhead.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Select TWO.)

**A.** Migrate remote users to AWS Client VPN

**B.** Migrate the home directories to Amazon FSx for Windows File Server.

**C.** Create an AWS Direct Connect connection from the on-premises data center to AWS.

**D.** Migrate the home directories to Amazon FSx for Lustre.

**E.** Create an AWS Storage Gateway Volume Gateway. Mount a volume from the Volume Gateway to the on-premises file server.

**Answer: A,B (LEAVE A REPLY)**

**NEW QUESTION: 131**

A company uses AWS Organizations to manage its development environment. Each development team at the company has its own AWS account Each account has a single VPC and CIDR blocks that do not overlap.

The company has an Amazon Aurora DB cluster in a shared services account All the development teams need to work with live data from the DB cluster Which solution will provide the required connectivity to the DB cluster with the LEAST operational overhead?

**A.** Create an AWS Resource Access Manager (AWS RAM) resource share tor the DB cluster. Share the DB cluster with all the development accounts

**B.** Create a transit gateway in the shared services account Create an AWS Resource Access Manager (AWS RAM) resource share for the transit gateway Share the transit gateway with all the development accounts Instruct the developers to accept the resource share Configure networking.

**C.** Create an Application Load Balancer (ALB) that points to the IP address of the DB cluster Create an AWS PrivateLink endpoint service that uses the ALB Add permissions to allow each development account to connect to the endpoint service

**D.** Create an AWS Site-to-Site VPN connection in the shared services account Configure networking Use AWS Marketplace VPN software in each development account to connect to the Site-to-Site VPN connection

**Answer: B (LEAVE A REPLY)**

* Create a Transit Gateway:
* In the shared services account, create a new AWS Transit Gateway. This serves as a central hub to connect multiple VPCs, simplifying the network topology and management.
* Configure Transit Gateway Attachments:
* Attach the VPC containing the Aurora DB cluster to the transit gateway. This allows the shared services VPC to communicate through the transit gateway.
* Create Resource Share with AWS RAM:
* Use AWS Resource Access Manager (AWS RAM) to create a resource share for the transit gateway. Share this resource with all development accounts. AWS RAM allows you to securely share your AWS resources across AWS accounts without needing to duplicate them.
* Accept Resource Shares in Development Accounts:
* Instruct each development team to log into their respective AWS accounts and accept the transit gateway resource share. This step is crucial for enabling cross-account access to the shared transit gateway.
* Configure VPC Attachments in Development Accounts:
* Each development account needs to attach their VPC to the shared transit gateway. This allows their VPCs to route traffic through the transit gateway to the Aurora DB cluster in the shared services account.
* Update Route Tables:
* Update the route tables in each VPC to direct traffic intended for the Aurora DB cluster through the transit gateway. This ensures that network traffic is properly routed between the development VPCs and the shared services VPC.

Using a transit gateway simplifies the network management and reduces operational overhead by providing a scalable and efficient way to interconnect multiple VPCs across different AWS accounts.

References

* AWS Database Blog on RDS Proxy for Cross-Account Access#48#.
* AWS Architecture Blog on Cross-Account and Cross-Region Aurora Setup#49#.

* DEV Community on Managing Multiple AWS Accounts with Organizations#51#.

**NEW QUESTION: 132**
An EC2-based ticketing service pulls a frequently updated pricing file (stored in S3) on startup. Sometimes EC2s have stale pricing, causing charge issues.
**A.** Lambda updates DynamoDB with new prices.
**B.** Lambda updates Amazon EFS.
**C.** Use Mountpoint for S3 to mount the pricing file to EC2.
**D.** Use Multi-Attach EBS volume for price file.
**Answer: C (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
Mountpoint for Amazon S3allows EC2 instances to directly access files in S3 as aPOSIX-compliant mount point, ensuring they always get the latest data without copying or syncing.
It's simple and cost-effective for read-heavy patterns.
#Mountpoint for Amazon S3

**NEW QUESTION: 133**
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

**Answer: B (LEAVE A REPLY)**
This option meets the RPO and RTO requirements for both the application and database tiers and uses tools like Amazon DLM and RDS automated backups to create and manage the backups. Additionally, it uses Global Accelerator to ensure low latency after failover by directing traffic to the closest healthy endpoint.

**NEW QUESTION: 134**
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
**Answer: B (LEAVE A REPLY)**
https://docs.aws.amazon.com/cur/latest/userguide/billing-cur-limits.html

**NEW QUESTION: 135**
Question:
A SaaS web app runs on EC2 Linux behind an ALB. It storesuser sessionsin an RDS Multi-AZ database. During high traffic, the app suffers latency due to session read/write.
What is the best way to reduce session latency?
Options:
**A.** Store session data in Amazon S3.
**B.** Use FSx for Windows and mount it.
**C.** Use Multi-Attach EBS volumes.
**D.** Use ElastiCache for Redis to store sessions.
**Answer: (SHOW ANSWER)**
* Dis the AWS best practice forsession storage: UseElastiCache for Redis- a fast, in-memory data store that handles high throughput with microsecond latency.
* It's highly scalable, fault-tolerant, and optimized for temporary, fast-access session data.
Incorrect:
* A: S3 is slow and object-based - not for session I/O.
* B: FSx is Windows-only and not ideal for this use case.
* C: EBS Multi-Attach has limitations, complexity, and is not suitable for high-performance shared memory.

Reference:#https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/WhatIs.html

**NEW QUESTION: 136**
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.
**Answer: (SHOW ANSWER)**
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.



**NEW QUESTION: 137**
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.
**Answer: (SHOW ANSWER)**
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.
https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html

**NEW QUESTION: 138**
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 AH 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 (Amazon CloudWatch Events) rule that runs once every day Configure the rule to invoke one AWS Lambda function that starts or stops instances based on the tag day and time.
**B.** Create an Amazon EventBridge (Amazon CloudWatch Events) 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 (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to invoke another Lambda function that starts instances based on the tag
**C.** Create an Amazon EventBridge (Amazon CloudWatch Events) rule that runs every business day in the evening Configure the rule to invoke an AWS Lambda function that terminates instances based on the tag Create a second EventBridge (CloudWatch Events) rule that runs every business day in the morning Configure the second rule to 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.
**Answer: B (LEAVE A REPLY)**
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.

**NEW QUESTION: 139**
A company runs an ecommerce web application on AWS. The static website is hosted on Amazon S3 and served via Amazon CloudFront. API Gateway invokes AWS Lambda for order processing, and Lambda stores data in an Amazon RDS for MySQL DB cluster (On-Demand Instances).
Recently, SQL injection attacks and latency during peak times (cold starts) have been reported. The company wants to ensure scalability, protect against web exploits, and reduce database costs.
**A.** Increase Lambda timeout, use RDS Reserved Instances, and use AWS Shield Advanced
**B.** Increase Lambda memory, switch to Redshift, use Amazon Inspector
**C.** Use provisioned concurrency, switch to Aurora Serverless, use AWS Shield Advanced
**D.** Use provisioned concurrency, use RDS Reserved Instances, use AWS WAF with CloudFront
**Answer: D (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
* Provisioned concurrencyaddresses cold starts and improves performance during spikes.
* RDS Reserved Instancessave money over On-Demand for predictable usage.
* AWS WAFis designed to protect againstSQL injectionandweb exploits, and integrates with CloudFront.
#AWS Lambda Provisioned Concurrency
#AWS WAF Web ACL with CloudFront

**NEW QUESTION: 140**
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 able in a new Region Use the new table as a replica target tor 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 able m 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 same Region, and configure S3 Same-Region Replication between the buckets- Create an Amazon CloudFront distribution and configure origin failover with two origin accessing the S3 buckets Create a new DynamoDB table m a new Region Use the new table as a replica target for DynamoDB global tables.
**Answer: C (LEAVE A REPLY)**
https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb-global-table-stream-lambda/?nc1=h_ls

**NEW QUESTION: 141**
A company is designing an AWS environment tor a manufacturing application. The application has been successful with customers, and the application's user base has increased. The company has connected the AWS environment to the company's on-premises data center through a 1 Gbps AWS Direct Connect connection. The company has configured BGP for the connection.
The company must update the existing network connectivity solution to ensure that the solution is highly available, fault tolerant, and secure.
Which solution win meet these requirements MOST cost-effectively?
**A.** Add a dynamic private IP AWS Site-to-Site VPN as a secondary path to secure data in transit and provide resilience for the Direct Conned connection. Configure MACsec to encrypt traffic inside the Direct Connect connection.
**B.** Provision another Direct Conned connection between the company's on-premises data center and AWS to increase the transfer speed and provide resilience. Configure MACsec to encrypt traffic inside the Dried Conned connection.
**C.** Configure multiple private VIFs. Load balance data across the VIFs between the on-premises data center and AWS to provide resilience.
**D.** Add a static AWS Site-to-Site VPN as a secondary path to secure data in transit and to provide resilience for the Direct Connect connection.

**Answer: A (LEAVE A REPLY)**
To enhance the network connectivity solution's availability, fault tolerance, and security in a cost-effective manner, adding a dynamic private IP AWS Site-to-Site VPN as a secondary path is a viable option. This VPN serves as a resilient backup for the Direct Connect connection, ensuring continuous data flow even if the primary path fails. Implementing MACsec (Media Access Control Security) on the Direct Connect connection further secures the data in transit by providing encryption, thus addressing the security requirement.
Thissolution strikes a balance between cost and operational efficiency, avoiding the higher expenses associated with provisioning an additional Direct Connect connection.
AWS Documentation on AWS Direct Connect and AWS Site-to-Site VPN provides insights into setting up resilient and secure network connections. Additionally, information on MACsec offers guidance on how to implement encryption for Direct Connect connections, aligning with best practices for secure and highly available network architectures.

**NEW QUESTION: 142**
A company is migrating an application from on-premises infrastructure to the AWS Cloud. During migration design meetings, the company expressed concerns about the availability and recovery options for its legacy Windows file server. The file server contains sensitive business-critical data that cannot be recreated in the event of data corruption or data loss. According to compliance requirements, the data must not travel across the public internet. The company wants to move to AWS managed services where possible.
The company decides to store the data in an Amazon FSx for Windows File Server file system. A solutions architect must design a solution that copies the data to another AWS Region for disaster recovery (DR) purposes.
Which solution will meet these requirements?
**A.** Create a destination Amazon S3 bucket in the DR Region. Establish connectivity between the FSx for Windows File Server file system in the primary Region and the S3 bucket in the DR Region by using Amazon FSx File Gateway. Configure the S3 bucket as a continuous backup source in FSx File Gateway.
**B.** Create an FSx for Windows File Server file system in the DR Region. Establish connectivity between the VPC in the primary Region and the VPC in the DR Region by using AWS Site-to-Site VPN. Configure AWS DataSync to communicate by using VPN endpoints.
**C.** Create an FSx for Windows File Server file system in the DR Region. Establish connectivity between the VPC in the primary Region and the VPC in the DR Region by using VPC peering. Configure AWS DataSync to communicate by using interface VPC endpoints with AWS PrivateLink.
**D.** Create an FSx for Windows File Server file system in the DR Region. Establish connectivity between the VPC in the primary Region and the VPC in the DR Region by using AWS Transit Gateway in each Region. Use AWS Transfer Family to copy files between the FSx for Windows File Server file system in the primary Region and the FSx for Windows File Server file system in the DR Region over the private AWS backbone network.
**Answer: (SHOW ANSWER)**

The best solution is to create an FSx for Windows File Server file system in the DR Region and establish connectivity between the VPCs in both Regions by using VPC peering. This will ensure that the data does not travel across the public internet and meets the compliance requirements. By using AWS DataSync with interface VPC endpoints and AWS PrivateLink, the data can be copied securely and efficiently between the FSx for Windows File Server file systems in both Regions. This solution also provides the ability to fail over to the DR Region in case of a disaster. References: [Amazon FSx for Windows File Server User Guide],
[AWS DataSync User Guide], [Amazon VPC User Guide]

**NEW QUESTION: 143**
A solutions architect has launched multiple Amazon EC2 instances in a placement group within a single Availability Zone. Because of additional load on the system, the solutions architect attempts to add new instances to the placement group. However, the solutions architect receives an insufficient capacity error. What should the solutions architect do to troubleshoot this issue?
**A.** Use a spread placement group. Set a minimum of eight instances for each Availability Zone.
**B.** Launch the additional instances as Dedicated Hosts in the placement groups.
**C.** Create a new placement group. Merge the new placement group with the original placement group.
**D.** Stop and start all the instances in the placement group. Try the launch again.
**Answer: D (LEAVE A REPLY)**

**NEW QUESTION: 144**
A company needs to migrate an on-premises SFTP site to AWS. The SFTP site currently runs on a Linux VM. Uploaded files are made available to downstream applications through an NFS share.
As part of the migration to AWS, a solutions architect must implement high availability. The solution must provide external vendors with a set of static public IP addresses that the vendors can allow. The company has set up an AWS Direct Connect connection between its on-premises data center and its VPC.
Which solution will meet these requirements with the least operational overhead?
**A.** Create an AWS Transfer Family server, configure an internet-facing VPC endpoint for the Transfer Family server, specify an Elastic IP address for each subnet, configure the Transfer Family server to pace files into an Amazon Elastic Files System (Amazon EFS) file system that is deployed across multiple Availability Zones Modify the configuration on the downstream applications that access the existing NFS share to mount the EFS endpoint instead.
**B.** Create an AWS Transfer Family server. Configure a publicly accessible endpoint for the Transfer Family server. Configure the Transfer Family server to place files into an Amazon Elastic Files System [Amazon EFS} the system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the its endpoint instead.
**C.** Use AWS Application Migration service to migrate the existing Linux VM to an Amazon EC2 instance. Assign an Elastic IP address to the EC2 instance. Mount an Amazon Elastic Fie system (Amazon EFS) the system to the EC2 instance. Configure the SFTP server to place files in. the EFS file system.

Modify the configuration on the downstream applications that access the existing NFS share to mount the EFS endpoint instead.

**D.** Use AWS Application Migration Service to migrate the existing Linux VM to an AWS Transfer Family server. Configure a publicly accessible endpoint for the Transfer Family server. Configure the Transfer Family sever to place files into an Amazon FSx for Luster the system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the FSx for Luster endpoint instead.

**Answer: (SHOW ANSWER)**

To migrate an on-premises SFTP site to AWS with high availability and a set of static public IP addresses for external vendors, the best solution is to create an AWS Transfer Family server with an internet-facing VPC endpoint. Assigning Elastic IP addresses to each subnet and configuring the server to store files in an Amazon Elastic File System (EFS) that spans multiple Availability Zones ensures high availability and consistent access. This approach minimizes operational overhead by leveraging AWS managed services and eliminates the need to manage underlying infrastructure.

AWS Documentation on AWS Transfer Family and Amazon Elastic File System provides detailed instructions on setting up a highly available SFTP environment on AWS. This solution is in line with AWS best practices for migrating and modernizing applications with minimal disruption and ensuring high availability and security.

**NEW QUESTION: 145**

A company needs to gather data from an experiment in a remote location that does not have internet connectivity. During the experiment, sensors that are connected to a total network will generate 6 TB of data in a preprimary formal over the course of 1 week. The sensors can be configured to upload their data files to an FTP server periodically, but the sensors do not have their own FTP server. The sensors also do not support other protocols. The company needs to collect the data centrally and move lie data to object storage in the AWS Cloud as soon. as possible after the experiment.

Which solution will meet these requirements?

**A.** Order an AWS Snowball Edge Compute Optimized device. Connect the device to the local network. Configure AWS DataSync with a target bucket name, and unload the data over NFS to the device. After the experiment return the device to AWS so that the data can be loaded into Amazon S3.

**B.** Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Create a shell script that periodically downloads data from each sensor. After the experiment, return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store [Amazon EBS) volume.

**C.** Order an AWS Snowcone device, including an Amazon Linux 2 AMI. Connect the device to the local network. Launch an Amazon EC2 instance on the device. Install and configure an FTP server on the EC2 instance. Configure the sensors to upload data to the EC2 instance. After the experiment, return the device to AWS so that the data can be loaded into Amazon S3.

**D.** Order an AWS Snowcone device. Connect the device to the local network. Configure the device to use Amazon FSx. Configure the sensors to upload data to the device. Configure AWS DataSync on the

device to synchronize the uploaded data with an Amazon S3 bucket Return the device to AWS so that the data can be loaded as an Amazon Elastic Block Store (Amazon EBS) volume.

**Answer: C (LEAVE A REPLY)**

For collecting data from remote sensors without internet connectivity, using an AWS Snowcone device with an Amazon EC2 instance running an FTP server presents a practical solution. This setup allows the sensors to upload data to the EC2 instance via FTP, and after the experiment, the Snowcone device can be returned to AWS for data ingestion into Amazon S3. This approach minimizes operational complexity and ensures efficient data transfer to AWS for further processing or storage.

AWS Documentation on AWS Snowcone and Amazon EC2 provides detailed guidance on deploying compute and storage capabilities in edge locations. This solution leverages AWS's edge computing devices to address challenges associated with data collection in remote or disconnected environments.

**NEW QUESTION: 146**

During an audit, a security team discovered that a development team was putting IAM user secret access keys in their code and then committing it to an AWS CodeCommit repository. The security team wants to automatically find and remediate instances of this security vulnerability.

Which solution will ensure that the credentials are appropriately secured automatically7

**A.** Run a script nightly using AWS Systems Manager Run Command to search tor credentials on the development instances. If found. use AWS Secrets Manager to rotate the credentials.

**B.** Use a scheduled AWS Lambda function to download and scan the application code from CodeCommit.

If credentials are found, generate new credentials and store them in AWS KMS.

**C.** Configure Amazon Made to scan for credentials in CodeCommit repositories. If credentials are found, trigger an AWS Lambda function to disable the credentials and notify the user.

**D.** Configure a CodeCommit trigger to invoke an AWS Lambda function to scan new code submissions for credentials. It credentials are found, disable them in AWS IAM and notify the user

**Answer: (SHOW ANSWER)**

CodeCommit may use S3 on the back end (and it also uses DynamoDB on the back end) but I don't think they're stored in buckets that you can see or point Macie to. In fact, there are even solutions out there describing how to copy your repo from CodeCommit into S3 to back it up:https://docs.aws.amazon.com /prescriptive-guidance/latest/patterns/automate-event-driven-backups-from-codecommit-to-amazon-s3-using- codebuild-and-cloudwatch-events.html

**NEW QUESTION: 147**

A company wants to send data from its on-premises systems to Amazon S3 buckets. The company created the S3 buckets in three different accounts. The company must send the data privately without the data traveling across the internet The company has no existing dedicated connectivity to AWS Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

**A.** Establish a networking account in the AWS Cloud Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a private VIF between the on-premises environment and the private VPC.

**B.** Establish a networking account in the AWS Cloud Create a private VPC in the networking account. Set up an AWS Direct Connect connection with a public VIF between the on-premises environment and the private VPC.
**C.** Create an Amazon S3 interface endpoint in the networking account.
**D.** Create an Amazon S3 gateway endpoint in the networking account.
**E.** Establish a networking account in the AWS Cloud Create a private VPC in thenetworking account. Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account.
**Answer: A,C (LEAVE A REPLY)**
https://docs.aws.amazon.com/AmazonS3/latest/userguide/privatelink-interface-endpoints.html#types-of-vpc- endpoints-for-s3
https://aws.amazon.com/premiumsupport/knowledge-center/s3-bucket-access-direct-connect/ Use a private IP address over Direct Connect (with an interface VPC endpoint) To access Amazon S3 using a private IP address over Direct Connect, perform the following steps:
3. Create a private virtual interface for your connection.
5. Create an interface VPC endpoint for Amazon S3 in a VPC that is associated with the virtual private gateway. The VGW must connect to a Direct Connect private virtual interface. This interface VPC endpoint resolves to a private IP address even if you enable a VPC endpoint for S3.

**NEW QUESTION: 148**
Question:
A company uses AWS Organizations and tags every resource with a BusinessUnit tag. They want toallocate cloud costsby business unit andvisualizethem.
Options:
**A.** Activate BusinessUnit cost allocation tag in the management account. Create a CUR to S3. Use Athena
+ QuickSight for reporting.
**B.** Create cost allocation tags in each member account. Use CloudWatch Dashboards.
**C.** Create cost allocation tags in the management account. Deploy CURs per account.
**D.** Use tags and CUR per account. Visualize with QuickSight from management account.
**Answer: A (LEAVE A REPLY)**
* Ais the correct best-practice approach:
* Activate cost allocation tags inmanagement/payer account.
* Enable AWS Cost and Usage Report (CUR) to dump usage to S3.
* Query CUR withAmazon Athena, and visualize withAmazon QuickSight.
Other options either duplicate CURs (C, D) or misuse tools (B: CloudWatch is not a cost analysis tool).
Reference:#https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-alloc-tags.html#https://docs.
aws.amazon.com/cur/latest/userguide/what-is-cur.html

**NEW QUESTION: 149**

A company completed a successful Amazon Workspaces proof of concept. They now want to make Workspaceshighly available across two AWS Regions. Workspaces are deployed in the failover Region. A hosted zone is available in Amazon Route 53.
What should the solutions architect do?
**A.** Create a connection alias in the primary Region and in the failover Region. Associate each with a directory in its Region. Create a Route 53 failover routing policy with Evaluate Target Health = Yes.
**B.** Create a connection alias in both Regions. Associate both with a directory in the primary Region. Use a Route 53 multivalue answer routing policy.
**C.** Create a connection alias in the primary Region. Associate with the directory in the primary Region. Use Route 53 weighted routing.
**D.** Create a connection alias in the primary Region. Associate it with the directory in the failover Region.Use Route 53 failover routing with Evaluate Target Health = Yes.
**Answer: (SHOW ANSWER)**
Comprehensive and Detailed Explanation:
A is correct because AWS recommends using oneconnection alias per Region, associated witheach directory. Then, configure a Route 53 failover policy so that if the primary Region becomes unhealthy, users are directed to the failover Region automatically. "Evaluate Target Health" ensures automatic detection and failover.
References:
Amazon Workspaces Cross-Region Resilience
Route 53 Failover Routing

**NEW QUESTION: 150**
A company is migrating an on-premises application and a MySQL database to AWS. The application processes highly sensitive data, and new data is constantly updated in the database. The data must not be transferred over the internet. The company also must encrypt the data in transit and at rest.
The database is 5 TB in size. The company already has created the database schema in an Amazon RDS for MySQL DB instance. The company has set up a 1 Gbps AWS Direct Connect connection to AWS. The company also has set up a public VIF and a private VIF. A solutions architect needs to design a solution that will migrate the data to AWS with the least possible downtime.
Which solution will meet these requirements?
**A.** Perform a database backup. Copy the backup files to an AWS Snowball Edge Storage Optimized device. Import the backup to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest. Use TLS for encryption in transit. Import the data from Amazon S3 to the DB instance.
**B.** Use AWS Database Migration Service (AWS DMS) to migrate the data to AWS. Create a DMS replication instance in a private subnet. Create VPC endpoints for AWS DMS. Configure a DMS task to copy data from the on-premises database to the DB instance by using full load plus change data capture (CDC). Use the AWS Key Management Service (AWS KMS) default key for encryption at rest. Use TLS for encryption in transit.

**C.** Perform a database backup. Use AWS DataSync to transfer the backup files to Amazon S3. Use server- side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest. Use TLS for encryption in transit. Import the data from Amazon S3 to the DB instance.
**D.** Use Amazon S3 File Gateway. Set up a private connection to Amazon S3 by using AWS PrivateLink.Perform a database backup. Copy the backup files to Amazon S3. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) for encryption at rest. Use TLS for encryption in transit. Import the data from Amazon S3 to the DB instance.
**Answer: B (LEAVE A REPLY)**
The best solution is to use AWS Database Migration Service (AWS DMS) to migrate the data to AWS. AWS DMS is a web service that can migrate data from various sources to various targets, including MySQL databases. AWS DMS can perform full load and change data capture (CDC) migrations, which means that it can copy the existing data and also capture the ongoing changes to keep the source and target databases in sync. This minimizes the downtime during the migration process. AWS DMS also supports encryption at rest and in transit by using AWS Key Management Service (AWS KMS) and TLS, respectively. This ensures that the data is protected during the migration. AWS DMS can also leverage AWS Direct Connect to transfer the data over a private connection, avoiding the internet. This solution meets all the requirements of the company. References: AWS Database Migration Service Documentation, Migrating Data to Amazon RDS for MySQL or MariaDB, Using SSL to Encrypt a Connection to a DB Instance

**NEW QUESTION: 151**
A financial company needs to create a separate AWS account for a new digital wallet application. The company uses AWS Organizations to manage its accounts. A solutions architect uses the 1AM user Supportl from the management account to create a new member account with finance1@example.com as the email address.
What should the solutions architect do to create IAM users in the new member account?
**A.** Sign in to the AWS Management Console with AWS account root user credentials by using the 64-character password from the initial AWS Organizations email senttofinance1@example.com. Set up the IAM users as required.
**B.** From the management account, switch roles to assume the OrganizationAccountAccessRole role with the account ID of the new member account. Set up the IAM users as required.
**C.** Go to the AWS Management Console sign-in page. Choose "Sign in using root account credentials." Sign in in by using the email address finance1@example.com and the management account's root password. Set up the IAM users as required.
**D.** Go to the AWS Management Console sign-in page. Sign in by using the account ID of the new member account and the Supportl IAM credentials. Set up the IAM users as required.
**Answer: D (LEAVE A REPLY)**
The best solution is to turn on the Concurrency Scaling feature for the Amazon Redshift cluster. This feature allows the cluster to automatically add additional capacity to handle bursts of read queries without affecting the performance of write queries. The additional capacity is transparent to the users and is billed separately based on the usage. This solution meets the business requirements of servicing read and

write queries at all times and is also cost-effective compared to the other options, which involve provisioning additional resources or resizing the cluster. References: Amazon Redshift Documentation, Concurrency Scaling in Amazon Redshift



**NEW QUESTION: 152**
A company is running a traditional web application on Amazon EC2 instances. The company needsto 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.
**Answer: B (LEAVE A REPLY)**
minimizes operational + microservices that run on containers = AWS Elastic Beanstalk

**NEW QUESTION: 153**
A company's solutions architect is reviewing a new internally developed application in a sandbox AWS account The application uses an AWS Auto Scaling group of Amazon EC2 instances that have an IAM instance profile attached Part of the application logic creates and accesses secrets from AWS Secrets Manager The company has an AWS Lambda function that calls the application API to test the

functionality The company also has created an AWS CloudTrail trail in the account The application's developer has attached the SecretsManagerReadWnte AWS managed IAM policy to an IAM role The IAM role is associated with the instance profile that is attached to the EC2 instances The solutions architect has invoked the Lambda function for testing The solutions architect must replace the SecretsManagerReadWnte policy with a new policy that provides least privilege access to the Secrets Manager actions that the application requires What is the MOST operationally efficient solution that meets these requirements?

**A.** Generate a policy based on CloudTrail events for the IAM role Use the generated policy output to create a new IAM policy Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role

**B.** Create an analyzer in AWS Identity and Access Management Access Analyzer Use the IAM role's Access Advisor findings to create a new IAM policy Use the newly created IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role

**C.** Use the aws cloudtrail lookup-events AWS CLI command to filter and export CloudTrail events that are related to Secrets Manager Use a new IAM policy that contains the actions from CloudTrailto replace the SecretsManagerReadWnte policy that is attached to the IAM role

**D.** Use the IAM policy simulator to generate an IAM policy for the IAM role Use the newly generated IAM policy to replace the SecretsManagerReadWnte policy that is attached to the IAM role

**Answer: B (LEAVE A REPLY)**

The IAM policy simulator will generate a policy that contains only the necessary permissions for the application to access Secrets Manager, providing the least privilege necessary to get the job done. This is the most efficient solution as it will not require additional steps such as analyzing CloudTrail events or manually creating and testing an IAM policy.

You can use the IAM policy simulator to generate an IAM policy for an IAM role by specifying the role and the API actions and resources that the application or service requires. The simulator will then generate an IAM policy that grants the least privilege access to those actions and resources.

Once you have generated an IAM policy using the simulator, you can replace the existing SecretsManagerReadWnte policy that is attached to the IAM role with the newly generated policy. This will ensure that the application or service has the least privilege access to the Secrets Manager actions that it requires.

You can access the IAM policy simulator through the IAM console, AWS CLI, and AWS SDKs. Here is the link for more information:https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_simulator.html

**NEW QUESTION: 154**

A company consists of 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 two- 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 encryptionkeys (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.
**Answer: A (LEAVE A REPLY)**
"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/

**NEW QUESTION: 155**
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
**Answer: B (LEAVE A REPLY)**
AWS Backup is a fully managed service that makes it easy to centralize and automate the creation, retention, and restoration of backups across AWS services. It provides a way to schedule automatic backups for CodeCommit repositories on an hourly basis. Additionally, it also supports cross-Region replication, which allows you to copy the backups to a second Region for disaster recovery.
By using AWS Backup, the company can set up an automatic and regular backup schedule for the CodeCommit repository, ensuring that the data is regularly backed up and stored in a second Region. This can provide a way to recover quickly from any disaster event that might occur.
Reference:
AWS Backup documentation:https://aws.amazon.com/backup/

AWS Backup for AWS CodeCommit documentation:https://aws.amazon.com/about-aws/whats-new/2020/07
/aws-backup-now-supports-aws-codecommit-repositories/

**NEW QUESTION: 156**
Question:
A company needs to migratesome Oracle databases to AWSwhile keeping otherson-premisesfor compliance. The on-prem databases containspatial dataand runcron jobs. The solution must allowquerying on-prem data as foreign tablesfrom AWS.
**A.** Use DynamoDB, SCT, and Lambda. Move spatial data to S3 and query with Athena.
**B.** Use RDS for SQL Server and AWS Glue crawlers for Oracle access.
**C.** Use EC2-hosted Oracle with Application Migration Service. Use Step Functions for cron.
**D.** Use RDS for PostgreSQL with DMS and SCT. Use PostgreSQL foreign data wrappers. Connectvia Direct Connect.
**Answer: D (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
D is correct becauseRDS for PostgreSQLsupportsforeign data wrappers (FDW)that allow querying remote Oracle databases. WithAWS Schema Conversion Tool (SCT)andDatabase Migration Service (DMS), schema and data can be migrated effectively.AWS Direct Connectensures secure, private connectivity to on- prem databases. Cron jobs can be run via EventBridge or external orchestration.
* A doesn't support relational/spatial querying.
* B doesn't support FDW or spatial types.
* C introduces unnecessary complexity.
#Reference:
https://www.postgresql.org/docs/current/postgres-fdw.html
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.PostgreSQL.html

**NEW QUESTION: 157**
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.** 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.

**C.** 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.
**D.** 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
**Answer: D (LEAVE A REPLY)**

**NEW QUESTION: 158**
A company migrated an application to the AWS Cloud. The application runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). Application data is stored in a MySQL database that runs on an additional EC2 instance. The application's use of the database is read-heavy.
The loads static content from Amazon Elastic Block Store (Amazon EBS) volumes that are attached to each EC2 instance. The static content is updated frequently and must be copied to each EBS volume.
The load on the application changes throughout the day. During peak hours, the application cannot handle all the incoming requests. Trace data shows that the database cannot handle the read load during peak hours.
Which solution will improve the reliability of the application?
**A.** Migrate the application to a set of AWS Lambda functions. Set the Lambda functions as targets for the ALB. Create a new single EBS volume for the static content. Configure the Lambda functions to read from the new EBS volume. Migrate the database to an Amazon RDS for MySQL Multi-AZ DB cluster.
**B.** Migrate the application to a set of AWS Step Functions state machines. Set the state machines as targets for the ALB. Create an Amazon Elastic File System (Amazon EFS) file system for the static content.
Configure the state machines to read from the EFS file system. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.
**C.** Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) Cluster. Use the AWS Fargate launch type for the tasks that host the application. Create a new single EBS volume the static content. Mount the new EBS volume on the ECS duster. Configure AWS Application Auto Scaling on ECS cluster. Set the ECS service as a target for the ALB. Migrate the database to an Amazon RDS for MySOL Multi-AZ DB cluster.
**D.** Containerize the application. Migrate the application to an Amazon Elastic Container Service (Amazon ECS) cluster. Use the AWS Fargate launch type for the tasks that host the application. Create an Amazon Elastic File System (Amazon EFS) file system for the static content. Mount the EFS file system to each container. Configure AWS Application Auto Scaling on the ECS cluster Set the ECS service as a target for the ALB. Migrate the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance.
**Answer: D (LEAVE A REPLY)**
This solution will improve the reliability of the application by addressing theissues of scalability, availability, and performance. Containerizing the application will make it easier to deploy and manage on AWS.

Migrating the application to an Amazon ECS cluster will allow the application to run on a fully managed container orchestration service. Using the AWS Fargate launch type for the tasks that host the application will enable the application to run on serverless compute engines that are automatically provisioned and scaled by AWS. Creating an Amazon EFS file system for the static content will provide a scalable and shared storage solution that can be accessed by multiple containers. Mounting the EFS file system to each container will eliminate the need to copy the static content to each EBS volume and ensure that the content is always up to date. Configuring AWS Application Auto Scaling on the ECS cluster will enable the application to scale up and down based on demand or a predefined schedule. Setting the ECS service as a target for the ALB will distribute the incoming requests across multiple tasks in the ECS cluster and improve the availability and fault tolerance of the application. Migrating the database to Amazon Aurora MySQL Serverless v2 with a reader DB instance will provide a fully managed, compatible, and scalable relational database service that can handle high throughput and concurrent connections. Using a reader DB instance will offload some of the read load from the primary DB instance and improve the performance of the database.

**NEW QUESTION: 159**
Question:
A company is replicating an application in asecondary Region. The application usesDynamoDBandRDS for MySQL. The secondary Region must function independently during adisaster.
**A.** Use DynamoDB global tables and an RDS read replica.
**B.** Use DAX and a read replica.
**C.** Use global tables and RDS Multi-AZ with standby in secondary Region.
**D.** Use Streams and Lambda to copy data. Use read replica.
**Answer: A (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
A is correct because:
* DynamoDB global tablesallow for multi-Region, active-active usage.
* RDS MySQL read replicain another Region supports read workloads and can bepromotedduring disaster to act as a standalone DB.
* B is incorrect: DAX is a cache, not a replication mechanism.
* C is wrong because Multi-AZ doesn't span Regions.
* D is more manual and error-prone.
#Reference:
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GlobalTables.html
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

**NEW QUESTION: 160**
A company is running an application in the AWS Cloud. The application uses AWS Lambda functions and Amazon Elastic Container Service (Amazon ECS) containers that run with AWS Fargate technology as its primary compute. The load on the application is irregular. The application experiences long periods of no usage, followed by sudden and significant increases and decreases in traffic. The application is write-

heavy and stores data in an Amazon Aurora MySQL database. The database runs on an Amazon RDS memory optimized DB instance that is not able to handle the load.
What is the MOST cost-effective way for the company to handle the sudden and significant changes in traffic?
**A.** Migrate the database to Aurora Serverless v1. Purchase Compute Savings Plans.
**B.** Add additional read replicas to the database. Purchase Instance Savings Plans and RDS Reserved Instances.
**C.** Migrate the database to an Aurora global database. Purchase Compute Savings Plans and RDS Reserved Instances.
**D.** Migrate the database to an Aurora multi-master DB cluster. Purchase Instance Savings Plans.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 161**
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.
**Answer: B (LEAVE A REPLY)**
Deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (Amazon ECR) and creating a Lambda layer that uses the Docker image as the source. Then, deploying the API's Lambda functions as Zip packages and configuring the packages to use the Lambda layer would meet the requirements for simplifying the deployment and optimizing for code reuse.
A Lambda layer is a distribution mechanism for libraries, custom runtimes, and other function dependencies.
It allows you to manage your in-development function code separately from your dependencies, this way you can easily update your dependencies without having to update your entire function code.

By deploying the shared libraries and custom classes to a Docker image and uploading the image to Amazon Elastic Container Registry (ECR), it makes it easy to manage and version the dependencies. This way, the company can use the same version of the dependencies across different Lambda functions.
By creating a Lambda layer that uses the Docker image as the source, the company can configure the API's Lambda functions to use the layer, reducing the need to include the dependencies in each function package, and making it easy to update the dependencies across all functions at once.
Reference:
AWS Lambda Layers documentation:https://docs.aws.amazon.com/lambda/latest/dg/configuration-layers.html AWS Elastic Container Registry (ECR) documentation:https://aws.amazon.com/ecr/ Building Lambda Layers with Docker documentation:https://aws.amazon.com/blogs/compute/building- lambda-layers-with-docker/

**NEW QUESTION: 162**
A company hosts its primary API on AWS using Amazon API Gateway and AWS Lambda functions. Internal applications and external customers use this API. Some customers also use a legacy API hosted on a standalone EC2 instance.
The company wants to increase security across all APIs to prevent denial of service (DoS) attacks, check for vulnerabilities, and guard against common exploits.
What should a solutions architect do to meet these requirements?
**A.** Use AWS WAF to protect both APIs. Configure Amazon Inspector to analyze the legacy API. Configure Amazon GuardDuty to monitor for malicious attempts to access the APIs.
**B.** Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to analyze both APIs. Configure Amazon GuardDuty to block malicious attempts.
**C.** Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to analyze the legacy API. Configure Amazon GuardDuty to monitor for malicious attempts to access the APIs.
**D.** Use AWS WAF to protect the API Gateway API. Configure Amazon Inspector to protect the legacy API. Configure Amazon GuardDuty to block malicious attempts.
**Answer: (SHOW ANSWER)**
Comprehensive and Detailed Explanation:
C is correct because:
* AWS WAF integrates natively with API Gateway and protects against common web exploits (e.g., SQL injection, XSS).
* Amazon Inspector can scan the legacy EC2 instance for known vulnerabilities.
* Amazon GuardDuty is a continuous security monitoring service that detects threats butdoes not block traffic (B and D are incorrect because GuardDuty doesn't block).
References:
AWS WAF Overview
Amazon Inspector Overview
Amazon GuardDuty Overview

**NEW QUESTION: 163**
A company has implemented a new security requirement According to the new requirement, the company must scan all traffic from corporate AWS instances in the company's VPC for violations of the company's security policies. As a result of these scans the company can block access to and from specific IP addresses.
To meet the new requirement, the company deploys a set of Amazon EC2 instances in private subnets to serve as transparent proxies The company installs approved proxy server software on these EC2 instances The company modifies the route tables on all subnets to use the corresponding EC2 instances with proxy software as the default route The company also creates security groups that are compliant with the security policies and assigns these security groups to the EC2 instances Despite these configurations, the traffic of the EC2 instances in their private subnets is not being properly forwarded to the internet.
What should a solutions architect do to resolve this issue?
**A.** Disable source'destination checks on the EC2 instances that run the proxy software
**B.** Add a rule to the security group that is assigned to the proxy EC2 instances to allow all traffic between instances that have this security group Assign this security group to all EC2 instances in the VPC.
**C.** Change the VPC's DHCP options set Set the DNS server options to point to the addresses of the proxy EC2 instances
**D.** Assign one additional elastic network interface to each proxy EC2 instance Ensure that one of these network interfaces has a route to the private subnets Ensure that the other network interface has a route to the internet.
**Answer: (SHOW ANSWER)**
* Identify Proxy EC2 Instances:
* Determine which EC2 instances in the private subnets are running the proxy server software.
* Disable Source/Destination Checks:
* For each of these EC2 instances, go to the AWS Management Console.
* Navigate to the EC2 dashboard, select the instance, and choose "Actions" > "Networking" > "Change Source/Dest.Check".
* Disable the source/destination check for these instances.
Disabling source/destination checks allows the EC2 instances to route traffic appropriately, enabling them to function as network appliances or proxies. This ensures that traffic from other instances in the private subnets can be routed through the proxy instances to the internet, meeting the company's security requirements.
References
* Amazon EC2 User Guide on Source/Destination Checks

**NEW QUESTION: 164**
A company is migrating an application to the AWS Cloud. The application runs in an on-premises data center and writes thousands of images into a mounted NFS file system each night. After the company migrates the application, the company will host the application on an Amazon EC2 instance with a mounted Amazon Elastic File System (Amazon EFS) file system.

The company has established an AWS Direct Connect connection to AWS. Before the migration cutover, a solutions architect must build a process that will replicate the newly created on-premises images to the EFS file system.
What is the MOST operationally efficient way to replicate the images?
**A.** Configure a periodic process to run the aws s3 sync command from the on-premises file system to Amazon S3. Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system.
**B.** Deploy an AWS Storage Gateway file gateway with an NFS mount point. Mount the file gateway file system on the on-premises server. Configure a process to periodically copy the images to the mount point.
**C.** Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an S3 bucket by using public VIF. Configure an AWS Lambda function to process event notifications from Amazon S3 and copy the images from Amazon S3 to the EFS file system.
**D.** Deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. Send data over the Direct Connect connection to an AWS PrivateLink int
**Answer: D (LEAVE A REPLY)**
This option uses AWS DataSync to replicate the on-premises images to the EFS file system over the Direct Connect connection. AWS DataSync is a service that automates and accelerates data transfer between on- premises storage systems and AWS storage services. It can transfer data to and from Amazon EFS, Amazon FSx for Windows File Server, and Amazon S3. To use AWS DataSync, the company needs to deploy an AWS DataSync agent to an on-premises server that has access to the NFS file system. The agent connects to the AWS DataSync service endpoint in the AWS Region where the EFS file system is located. The company can use an AWS PrivateLink interface endpoint to connect to the service endpoint securely and privately over the Direct Connect connection. The company can then create a task in AWS DataSync that specifies the source location (the NFS file system), the destination location (the EFS file system), and the options for the data transfer (such as schedule, bandwidth limit, and verification). AWS DataSync will then perform the data transfer efficiently and securely, using encryption in transit and at rest.

**NEW QUESTION: 165**
A company has several AWS accounts. A development team is building an automation framework for cloud governance and remediation processes. The automation framework uses AWS Lambda functions in a centralized account. A solutions architect must implement a least privilege permissions policy that allows the Lambda functions to run in each of the company's AWS accounts.
Which combination of steps will meet these requirements? (Choose two.)
**A.** In the centralized account, create an IAM role that has the Lambda service as a trusted entity. Add an inline policy to assume the roles of the other AWS accounts.
**B.** In the other AWS accounts, create an IAM role that has minimal permissions. Add the centralized account's Lambda IAM role as a trusted entity.
**C.** In the centralized account, create an IAM role that has roles of the other accounts as trusted entities.

Provide minimal permissions.
**D.** In the other AWS accounts, create an IAM role that has permissions to assume the role of the centralized account. Add the Lambda service as a trusted entity.
**E.** In the other AWS accounts, create an IAM role that has minimal permissions. Add the Lambda service as a trusted entity.
**Answer: A,B (LEAVE A REPLY)**
https://medium.com/@it.melnichenko/invoke-a-lambda-across-multiple-aws-accounts-8c094b2e70be

**NEW QUESTION: 166**
Question:
A company runs a Linux app on Amazon EKS usingM6iEC2 instances under a Savings Plan that is about to expire. They want toreduce costsafter expiration.
**A.** Rebuild containers forARM64architecture.
**B.** Rebuild containers for container compatibility (invalid/unclear).
**C.** Migrate EKS nodes toGraviton(e.g., C7g, M7g).
**D.** Replace nodes with latestx86_64instances.
**E.** Purchase new Savings Plan for Graviton instance family.
**F.** Purchase new Savings Plan for x86_64 instances.
**Answer: A,C,E (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
* A: To run on Graviton, containers must supportARM64.
* C: Graviton-based EC2 instances offer significant cost savings and better price-performance.
* E: Once migrated, aSavings Planfor the new instance family ensures additional cost optimization.
* B is a non-sensical option.
* D and F continue with x86, which is more expensive.
#Reference:AWS Graviton Instances



**NEW QUESTION: 167**
A company has an application that uses an Amazon Aurora PostgreSQL DB cluster for the application's database. The DB cluster contains one small primary instance and three larger replica instances. The application runs on an AWS Lambda function. The application makes many short-lived connections to the database's replica instances to perform read-only operations.

During periods of high traffic, the application becomes unreliable and the database reports that too many connections are being established. The frequency of high-traffic periods is unpredictable.
Which solution will improve the reliability of the application?
**A.** Increase the max_connections setting on the DB cluster's parameter group. Reboot all the instances in the DB cluster. Update the Lambda function to connect to the DB cluster endpoint.
**B.** Use Amazon RDS Proxy to create a proxy for the DB cluster. Configure a read-only endpoint for the proxy. Update the Lambda function to connect tothe proxyendpoint.
**C.** Configure instance scaling for the DB cluster to occur when the DatabaseConnections metric is close to the max _ connections setting. Update the Lambda function to connect to the Aurora reader endpoint.
**D.** Use Amazon RDS Proxy to create a proxy for the DB cluster. Configure a read-only endpoint for the Aurora Data API on the proxy. Update the Lambda function to connect to the proxy endpoint.
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 168**
A company needs to store and process image data that will be uploaded from mobile devices using a custom mobile app. Usage peaks between 8 AM and 5 PM on weekdays, with thousands of uploads per minute. The app is rarely used at any other time. A user is notified when image processing is complete.
Which combination of actions should a solutions architect take to ensure image processing can scale to handle the load? (Select THREE.)
**A.** Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon MQ queue.
**B.** Upload files from the mobile software directly to Amazon S3. Use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SOS) standard queue.
**C.** Invoke an AWS Lambda function to perform image processing when a message is available in the queue.
**D.** Invoke an S3 Batch Operations job to perform image processing when a message is available in the queue
**E.** Send a push notification to the mobile app by using Amazon Simple Notification Service (Amazon SNS) when processing is complete.
**F.** Send a push notification to the mobile app by using Amazon Simple Email Service (Amazon SES) when processing is complete.
**Answer: (SHOW ANSWER)**
The best solution is to upload files from the mobile software directly to Amazon S3, use S3 event notifications to create a message in an Amazon Simple Queue Service (Amazon SQS) standard queue, and invoke an AWS Lambda function to perform image processing when a message is available in the queue. This solution will ensure that image processing can scale to handle the load, as Amazon S3 can store any amount of data and handle concurrent uploads, Amazon SQS can buffer the messages and deliver them reliably, and AWS Lambda can run code without provisioning or managing servers and scale automatically based on the demand. This solution will also notify the user when processing is complete by sending a push notification to the mobile app using Amazon Simple Notification Service (Amazon SNS), which is a web service that enables applications to send and receive notifications from the cloud.

This solution is more cost-effective than using Amazon MQ, which is a managed message broker service for Apache ActiveMQ that requires a dedicated broker instance, or S3 Batch Operations, which is a feature that allows users to perform bulk actions on S3 objects, such as copying or tagging, but does not support custom code execution. This solution is also more suitable than using Amazon Simple Email Service (Amazon SES), which is a web service that enables applications to send and receive email messages, but does not support push notifications for mobile devices. References: Amazon S3 Documentation, Amazon SQS Documentation, AWS Lambda Documentation, Amazon SNS Documentation

**NEW QUESTION: 169**

A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC. and the instances have access to the internet.

The company's security policy requires that partners can access resources only from domains that the company owns.

Which solution will meet these requirements?

**A.** Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. Configure a rule that has a low numeric value that allows requests for domains in the allowed list. Associate the rule group with the VPC.

**B.** Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a Route 53 outbound endpoint. Associate the outbound endpoint with the VPC. Associate the domain list with the outbound endpoint.

**C.** Create an Amazon Route 53 traffic flow policy to match the allowed domains. Configure the traffic flow policy to forward requests that match to the Route 53 Resolver. Associate the traffic flow policy with the VPC.

**D.** Create an Amazon Route 53 outbound endpoint. Associate the outbound endpoint with the VPC.Configure a Route 53 traffic flow policy to forward requests for allowed domains to the outbound endpoint. Associate the traffic flow policy with the VPC.

**Answer: A (LEAVE A REPLY)**

The company should create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. The company should configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. The company should configure a rule that has a low numeric value that allows requests for domains in the allowed list. The company should associate the rule group with the VPC.

This solution will meet the requirements because Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can createreusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs. You can specify lists of domain names to allow or block, and you can customize theresponses for the DNS queries that you block1.

By creating a domain list with the allowed domains and a rule group with rules to allow or block requests based on the domain list, the company can enforce its security policy and control access to sites.
The other options are not correct because:
* Configuring a Route 53 outbound endpoint and associating it with the VPC would not help with filtering outbound DNS traffic. A Route 53 outbound endpoint is a resource that enables you to forward DNS queries from your VPC to your network over AWS Direct Connect or VPN connections2.It does not provide any filtering capabilities.
* Creating a Route 53 traffic flow policy to match the allowed domains would not help with filtering outbound DNS traffic. A Route 53 traffic flow policy is a resource that enables you to route traffic based on multiple criteria, such as endpoint health, geographic location, and latency3.It does not provide any filtering capabilities.
* Creating a Gateway Load Balancer (GWLB) would not help with filtering outbound DNS traffic. A GWLB is a service that enables you to deploy, scale, and manage third-party virtual appliances such as firewalls, intrusion detection and prevention systems, and deep packet inspection systems in the cloud4.
It does not provide any filtering capabilities.
References:
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-outbound-endpoints.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/traffic-flow.html
https://docs.aws.amazon.com/elasticloadbalancing/latest/gateway/introduction.html

**NEW QUESTION: 170**
A company is collecting a large amount of data from a fleet of IoT devices Data is stored as Optimized Row Columnar (ORC) files in the Hadoop Distributed File System (HDFS) on a persistent Amazon EMR cluster.
The company's data analytics team queries the data by using SQL in Apache Presto deployed on the same EMR cluster Queries scan large amounts of data, always run for less than 15 minutes, and run only between 5 PM and 10 PM.
The company is concerned about the high cost associated with the current solution A solutions architect must propose the most cost-effective solution that will allow SQL data queries Which solution will meet these requirements?
**A.** Store data in Amazon S3 Use Amazon Redshift Spectrum to query data.
**B.** Store data in Amazon S3 Use the AWS Glue Data Catalog and Amazon Athena to query data
**C.** Store data in EMR File System (EMRFS) Use Presto in Amazon EMR to query data
**D.** Store data in Amazon Redshift. Use Amazon Redshift to query data.
**Answer: (SHOW ANSWER)**
(https://stackoverflow.com/questions/50250114/athena-vs-redshift-spectrum)

**NEW QUESTION: 171**

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 VPC. Create 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.
**Answer: C (LEAVE A REPLY)**
This solution requires the least amount of effort as it only requires to update the API endpoint to private in API Gateway and create an interface VPC endpoint. Then create a resource policy and attach it to the API.
This will make the API only accessible from the VPC and still keep the authentication mechanism intact.
Reference:
* https://aws.amazon.com/premiumsupport/knowledge-center/private-api-gateway-vpc-endpoint/
* https://aws.amazon.com/api-gateway/features/

**NEW QUESTION: 172**
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 ail 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 thatbelong to the global offices Developers Will reference this list to gain access to 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 involved when the JSON file is updated. Subscribe an AWS Lambda function to the SNS topic to update all relevant security group rules with Vie 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 is 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 me other accounts to reference the transit account's security group by using a nested security group reference of *<transit-account-id>./sg-1a2b3c4d".
**Answer: (SHOW ANSWER)**
Customer-managed prefix lists - Sets of IP address ranges that you define and manage. You can share your prefix list with other AWS accounts, enabling those accounts to reference the prefix list in their own resources.
https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html 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.

**NEW QUESTION: 173**
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 win deploy a two-net application In a new VPC To simplify the configuration, the db. example com CNAME record set tor 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? (Select 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 /eto /resolv.conf file
**C.** Create an authorization lo 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 m 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.
**Answer: (SHOW ANSWER)**
https://aws.amazon.com/premiumsupport/knowledge-center/private-hosted-zone-different-account/

**NEW QUESTION: 174**

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 VPC. Create a Route 53 outbound 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 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.
**Answer: (SHOW ANSWER)**
Amazon Route 53 Resolver is a managed DNS resolver service from Route 53 that helps to create conditional forwarding rules to redirect query traffic1. By associating the private hosted zone to all the VPCs, the solutions architect can enable DNS resolution for cloud.example.com within the VPCs. By creating a Route
53 inbound resolver in the shared services VPC, the solutions architect can enable DNS resolution for cloud.
example.com from on-premises systems. By attaching all VPCs to the transit gateway, the solutions architect can enable connectivity between the VPCs and the on-premises network through AWS Direct Connect. By creating forwarding rules in the on-premises DNS server for cloud.example.com that point to the inbound resolver, the solutions architect can direct DNS queries for cloud.example.com to the Route 53 Resolver endpoint in AWS. This solution will provide the highest performance as it leverages Route 53 Resolver's optimized routing and caching capabilities.
References: 1: https://aws.amazon.com/route53/resolver/

**NEW QUESTION: 175**
An adventure company has launched a new feature on its mobile app. Users can use the feature to upload their hiking and ratting 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.
**Answer: A (LEAVE A REPLY)**
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.

**NEW QUESTION: 176**
A company needs to create and manage multiple AWS accounts for a number of departments from a central location. The security team requires read-only access to all accounts from its own AWS account. The company is using AWS Organizations and created an account for the security team.
How should a solutions architect meet these requirements?
**A.** Use the OrganizationAccountAccessRole IAM role to create a new IAM policy with read-only access in each member account. Establish a trust relationship between the IAM policy in each member account and the security account. Ask the security team to use the IAM policy to gain access.
**B.** Use the Organization AccountAccessRole IAM role to create a new IAM role with read-only access in each member account. Establish a trust relationship between the IAM role in each member account and the security account. Ask the security team to use the IAM role to gain access.
**C.** Ask the security team to use AWS Security Token Service (AWS STS) lo call the AssumeRole API tor the Organization AccountAccessRole IAM role in the management account from the security account. Use the generated temporary credentials to gain access.
**D.** Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the Organization AccountAccessRole IAM role in the member account from the security account. Use the generated temporary credentials to gain access.
**Answer: B (LEAVE A REPLY)**
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html#orgs_manage_accounts_access-cross-account-role "When you create a member account using the AWS Organizations console, AWS Organizations automatically creates an IAM role named

OrganizationAccountAccessRole in the account" you need OrganizationAccountAccessRole in member account to create an read-only role and use role from security team to assume this read-only role.

**NEW QUESTION: 177**
Question:
A company runs workloads on EC2 inmultiple VPCsin a single Region. They also have anon-premises DNS server(via Direct Connect). All EC2 instances must resolve internal.company.com usingprivate communication.
What should a solutions architect do? (Select THREE.)
Options:
**A.** Create an Amazon Route 53inbound endpointin all workload VPCs.
**B.** Create a Route 53outbound endpointin one VPC.
**C.** Create a Route 53forwarding ruleto forward internal.company.com to the on-prem DNS.
**D.** Create a Route 53 rule with theSystemtype.
**E.** Associate the rule with all VPCs.
**F.** Associate the rule only with the VPC that has the outbound endpoint.
**Answer: (SHOW ANSWER)**
* B: Outbound Resolver endpoints enable EC2 instances tosendDNS queries to external resolvers (i.e., on-prem DNS servers).
* C: A forwarding rule must be created to forward the internal.company.com domain to the on-prem resolver.
* E: You must associate the rule withevery VPCthat needs to forward the queries.
Incorrect:
* A: Inbound endpoints are forreceivingDNS queries from on-prem to AWS, not vice versa.
* D: There's no such thing as a "System" rule type.
* F: Limiting to one VPC won't meet the requirement.
Reference:#https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html

**NEW QUESTION: 178**
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? (Select THREE.)
**A.** Activate the user-defined 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.
**Answer: A,C,F (LEAVE A REPLY)**

https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.html
https://aws.amazon.com/premiumsupport/knowledge-center/cost-explorer-analyze-spending-and-usage/
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/manage-cost-categories.htmlhttps://docs.aws.amazon.com/cost-management/latest/userguide/ce-enable.html
The best combination of actions to meet the company's requirements is Options A, C, and F.
Option A involves activating the user-defined cost allocation tags that represent the application and the team.
This will allow the company to assign costs to different applications or teams, and will allow them to be tracked in the monthly AWS bill.
Option C involves creating a cost category for each application in Billing and Cost Management. This will allow the company to easily identify and compare costs across different applications and teams.
Option F involves enabling Cost Explorer. This will allow the company to view the costs of their AWS resources over the last 12 months and to create forecasts for the next 12 months.
These recommendations are in line with the official Amazon Textbook and Resources for the AWS Certified Solutions Architect - Professional certification. In particular, the book states that "You can use cost allocation tags to group your costs by application, team, or other categories" (Source:https://d1.awsstatic.com/training- and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf).
Additionally, the book states that "Cost Explorer enables you to view the costs of your AWS resources over the last 12 months and to create forecasts for the next 12 months" (Source:https://d1.awsstatic.com/training- and-certification/docs-sa-pro/AWS_Certified_Solutions_Architect_Professional_Exam_Guide_EN_v1.5.pdf).

**NEW QUESTION: 179**
Question:
A company runs production workloads on EC2 On-Demand Instances and RDS for PostgreSQL. They want to reduce costs without compromising availability or capacity.
**A.** Use CUR and Lambda to terminate underutilized instances. Buy Savings Plans.
**B.** Use Budgets and Trusted Advisor, then manually terminate and buy RIs.
**C.** UseCompute OptimizerandTrusted Advisorfor recommendations. Apply rightsizing, auto scaling, and purchase a Compute Savings Plan.
**D.** Use Cost Explorer, alerts, and replace with Spot Instances.
**Answer: C (LEAVE A REPLY)**
Comprehensive and Detailed Explanation:
C is correct:AWS Compute Optimizeruses machine learning to analyze usage patterns and recommends rightsizing.Trusted Advisoradds further insights. Combining these withSavings Plansgives the best cost optimizationwithout reducing availability.
* A is risky due to using Lambda for termination.
* B and D offer partial or manual solutions.
#Reference:AWS Compute Optimizer

**NEW QUESTION: 180**
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.
**Answer: A,C,D (LEAVE A REPLY)**
The company wants to serve static content from an S3 bucket during the maintenance period. To do this, the following steps are required:
* Upload static informational content to the S3 bucket. This will provide the source of the content that will be served to the visitors.
* 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). This will allow CloudFront to access the S3 bucket securely and prevent public access to the bucket.
* During the weekly maintenance, edit the default cache behavior to use the S3 origin. Revert the change when the maintenance is complete. This will redirect all web requests to the S3 bucket instead of the Elastic Beanstalk domain name.
The other options are not correct because:
* Creating a new CloudFront distribution is not necessary and would require changing the alternate domain name configuration.
* Creating a cache behavior for the S3 origin on a new distribution would not work because the visitors would still access the original distribution using the alternate domain name.
* Configuring Elastic Beanstalk to serve traffic from the S3 bucket is not possible and would not achieve the desired result.
References:

https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide
/DownloadDistS3AndCustomOrigins.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to- s3.html
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/distribution-web-values-specify.html#DownloadDistValuesPathPattern

**NEW QUESTION: 181**
A company is using AWS Organizations to manage multiple accounts Due to regulatory requirements, the company wants to restrict specific member accounts to certain AWS Regions, where they are permitted to deploy resources The resources in the accounts must be tagged enforced based on a group standard and centrally managed with minimal configuration.
What should a solutions architect do to meet these requirements'?
**A.** Create an AWS Config rule in the specific member accounts to limit Regions and apply a tag policy.
**B.** From the AWS Billing and Cost Management console in the management account, disable Regions for the specific member accounts and apply a tag policy on the root.
**C.** Associate the specific member accounts with the root Apply a tag policy and an SCP using conditions to limit Regions.
**D.** Associate the specific member accounts with a new OU. Apply a tag policy and an SCP using conditions to limit Regions.
**Answer: D (LEAVE A REPLY)**
https://aws.amazon.com/es/blogs/mt/implement-aws-resource-tagging-strategy-using-aws-tag-policies-and- service-control-policies-scps/



**NEW QUESTION: 182**
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.
**Answer: A (LEAVE A REPLY)**
https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/custom-tags.htmlhttps://docs.aws.amazon.com
/awsaccountbilling/latest/aboutv2/configurecostallocreport.html

**NEW QUESTION: 183**
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 csvfile to an Amazon QuickSight dashboard.
**Answer: B (LEAVE A REPLY)**
https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage_lens.html

**NEW QUESTION: 184**

A company is building an application that will run on an AWS Lambda function. Hundreds of customers will use the application. The company wants to give each customer a quota of requests for a specific time period.
The quotas must match customer usage patterns. Some customers must receive a higher quota for a shorter time period.
Which solution will meet these requirements?
**A.** Create an Amazon API Gateway REST API with a proxy integration to invoke the Lambda function. For each customer, configure an API Gateway usage plan that includes an appropriate request quota. Create an API key from the usage plan for each user that the customer needs.
**B.** Create an Amazon API Gateway HTTP API with a proxy integration to invoke the Lambda function. For each customer, configure an API Gateway usage plan that includes an appropriate request quota. Configure route-level throttling for each usage plan. Create an API key from the usage plan for each user that the customer needs.
**C.** Create a Lambda function alias for each customer. Include a concurrency limit with an appropriate request quota. Create a Lambda function URL for each function alias. Share the Lambda function URL for each alias with the relevant customer.
**D.** Create an Application Load Balancer (ALB) in a VPC. Configure the Lambda function as a target for the ALB. Configure an AWS WAF web ACL for the ALB. For each customer, configure a rate-based rule that includes an appropriate request quota.
**Answer: A (LEAVE A REPLY)**
The correct answer is A.
A: This solution meets the requirements because it allows the company to create different usage plans for each customer, with different request quotas and time periods. The usage plans can be associated with API keys, which can be distributed to the users of each customer. The API Gateway REST API can invoke the Lambda function using a proxy integration, which passes the request data to the function as input and returns the function output as the response.This solution is scalable, secure, and cost-effective12 B: This solution is incorrect because API Gateway HTTP APIs do not support usage plans or API keys.These features are only available for REST APIs3 C: This solution is incorrect because it does not provide a way to enforce request quotas for each customer.
Lambda function aliases can be used to create different versions of the function, but they do not have any quota mechanism.Moreover, this solutionexposes the Lambda function URLs directly to the customers, which is not secure or recommended4 D: This solution is incorrect because it does not provide a way to differentiate between customers or users.
AWS WAF rate-based rules can be used to limit requests based on IP addresses, but they do not support any other criteria such as user agents or headers.Moreover, this solution adds unnecessary complexity and cost by using an ALB and a VPC56 References:
1:Creating and using usage plans with API keys - Amazon API Gateway2:Set up a proxy integration with a Lambda proxy integration - Amazon API Gateway3:Choose between HTTP APIs and REST APIs - Amazon API Gateway4:Using AWS Lambdaaliases - AWS Lambda5:Rate-based rule statement - AWS WAF, AWS Firewall Manager, and AWS Shield Advanced6:Lambda functions as targets for Application Load Balancers - Elastic Load Balancing

**NEW QUESTION: 185**
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.** 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.
**B.** 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 password rotation.
**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 RDS for Oracle. Store the password in AWS Secrets Manager. Turn on automatic rotation. Configure a yearly rotation schedule.
**Answer: D (LEAVE A REPLY)**

**NEW QUESTION: 186**
A company that develops consumer electronics with offices in Europe and Asia has 60 TB of software images stored on premises in Europe The company wants to transfer the images to an Amazon S3 bucket in the ap- northeast-1 Region New software images are created daily and must be encrypted in transit The company needs a solution that does not require custom development to automatically transfer all existing and new software images to Amazon S3 What is the next step in the transfer process?
**A.** Deploy an AWS DataSync agent and configure a task to transfer the images to the S3 bucket
**B.** Configure Amazon Kinesis Data Firehose to transfer the images using S3 Transfer Acceleration
**C.** Use an AWS Snowball device to transfer the images with the S3 bucket as the target
**D.** Transfer the images over a Site-to-Site VPN connection using the S3 API with multipart upload
**Answer: A (LEAVE A REPLY)**
* Deploy AWS DataSync Agent:
* Install the DataSync agent on your on-premises environment. This can be done by downloading the agent as a virtual appliance and deploying it on VMware ESXi, Hyper-V, or KVM hypervisors.
* Configure Source and Destination Locations:
* Set up the source location pointing to your on-premises storage where the software images are currently stored.
* Configure the destination location to point to your Amazon S3 bucket in the ap-northeast-1 Region.
* Create and Schedule DataSync Tasks:
* Create a DataSync task to automate the transfer process. This task will specify the source and destination locations and set options for how the data should be transferred.

* Schedule the task to run at intervals that suit your data transfer requirements, ensuring new images are transferred as they are created.
* Encryption in Transit:
* AWS DataSync automatically encrypts data in transit using TLS, ensuring that your data is secure during the transfer process.
* Monitoring and Management:
* Use the DataSync console or the AWS CLI to monitor the progress of your data transfers and manage the tasks.
AWS DataSync is an efficient solution that automates and accelerates the process of transferring large amounts of data to AWS, handling encryption, data integrity checks, and optimizing network usage without requiring custom development.
References
* AWS Storage Blog on DataSync#40#.
* AWS DataSync Documentation#41#.

**NEW QUESTION: 187**
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.
**Answer: (SHOW ANSWER)**
https://docs.aws.amazon.com/autoscaling/ec2/userguide/create-asg-instance-type-requirements.html#use- attribute-based-instance-type-selection-prerequisites

**NEW QUESTION: 188**
A company has a legacy application that runs on multiple .NET Framework components. The components share the same Microsoft SQL Server database and communicate with each other asynchronously by using Microsoft Message Queueing (MSMQ).
The company is starting a migration to containerized .NET Core components and wants to refactor the application to run on AWS. The .NET Core components require complex orchestration. The company must have full control over networking and host configuration. The application's database model is stronglyrelational.
Which solution will meet these requirements?

**A.** Host the .NET Core components on AWS App Runner. Host the database on Amazon RDS for SQL Server. Use Amazon EventBridge for asynchronous messaging.
**B.** Host the .NET Core components on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Host the database on Amazon DynamoDB. Use Amazon Simple Notification Service (Amazon SNS) for asynchronous messaging.
**C.** Host the .NET Core components on AWS Elastic Beanstalk. Host the database on Amazon Aurora PostgreSQL Serverless v2. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) for asynchronous messaging.
**D.** Host the .NET Core components on Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Host the database on Amazon Aurora MySQL Serverless v2. Use Amazon Simple Queue Service (Amazon SQS) for asynchronous messaging.
**Answer: (SHOW ANSWER)**
Hosting the .NET Core components on Amazon ECS with the Amazon EC2 launch type will meet the requirements of having complex orchestration and full control over networking and host configuration. Amazon ECS is a fully managed container orchestration service that supports both AWS Fargate and Amazon EC2 as launch types. The Amazon EC2 launch type allows users to choose their own EC2 instances, configure their own networking settings, and access their own host operating systems. Hosting the database on Amazon Aurora MySQL Serverless v2 will meet the requirements of having a strongly relational database model and using the same database engine as SQL Server. MySQL is a compatible relational database engine with SQL Server, and it can support most of the legacy application's database model. Amazon Aurora MySQL Serverless v2 is a serverless version of Amazon Aurora MySQL that can scale up and down automatically based on demand. Using Amazon SQS for asynchronous messaging will meet the requirements of providing a compatible replacement for MSMQ, which is a queue-based messaging system3. Amazon SQS is a fully managed message queuing service that enables decoupled and scalable microservices, distributed systems, and serverless applications.

**NEW QUESTION: 189**
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 tomerge 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.
**Answer: D (LEAVE A REPLY)**
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.

**NEW QUESTION: 190**
A Solutions Architect wants to make sure that only AWS users or roles with suitable permissions can access a new Amazon API Gateway endpoint. The Solutions Architect wants an end-to-end view of each request to analyze the latency of the request and create service maps.
How can the Solutions Architect design the API Gateway access control and perform request inspections?
**A.** For the API Gateway method, set the authorization to AWS_IAM. Then, give the IAM user or role execute-api:Invoke permission on the REST API resource. Enable the API caller to sign requests with AWS Signature when accessing the endpoint. Use AWS X-Ray to trace and analyze user requests to API Gateway.
**B.** For the API Gateway resource, set CORS to enabled and only return the company's domain in Access- Control-Allow-Origin headers. Then, give the IAM user or role execute-api:Invoke permission on the REST API resource. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.
**C.** Create a client certificate for API Gateway. Distribute the certificate to the AWS users and roles that need to access the endpoint. Enable the API caller to pass the client certificate when accessing the endpoint. Use Amazon CloudWatch to trace and analyze user requests to API Gateway.
**D.** Create an AWS Lambda function as the custom authorizer, ask the API client to pass the key and secret when making the call, and then use Lambda to validate the key/secret pair against the IAM system. Use AWS X-Ray to trace and analyze user requests to API Gateway.
**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 191**

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.
**Answer: B,D (LEAVE A REPLY)**
Moving the application frontend to a static website that is hosted on Amazon S3 will save cost as S3 is cheaper than running EC2 instances.
Using Spot instances for the backend EC2 instances will also save cost, as they are significantly cheaper than On-Demand instances. This will be suitable for the application, as it has minimal traffic during the rest of the day, and the availability of spot instances will not negatively affect the application's availability.
Reference:
Amazon S3 pricing:https://aws.amazon.com/s3/pricing/
Amazon EC2 Spot Instances documentation:https://aws.amazon.com/ec2/spot/ AWS Elastic Beanstalk documentation:https://aws.amazon.com/elasticbeanstalk/ Amazon Elastic Compute Cloud (EC2) pricing:https://aws.amazon.com/ec2/pricing/

**NEW QUESTION: 192**
Question:
A company uses IAM Identity Center for data scientist access. Each user should be able to accessonly their own datain an S3 bucket. The company also needs to generatemonthly access reportsper user.
Options:
**A.** Use IAM Identity Center permission sets to allow S3 access scoped to userName tag.
**B.** Use a shared IAM Identity Center role for all users and bucket policy.
**C.** Use AWS CloudTrail to log S3 data events, query via Athena.
**D.** Use CloudTrail management events to CloudWatch, then use Athena.
**E.** Use S3 access logs and S3 Select for reporting.
**Answer: A,C (LEAVE A REPLY)**

* A: Use dynamic IAM policies with {aws:PrincipalTag/userName} to enforceprefix-level access control - i.e., bucket/userA/*, bucket/userB/*.
* C: Enable CloudTraildata eventsto capture object-level access andquery them withAthena. This is the AWS-recommended way to audit per-user object access.
Incorrect:
* B doesn't provide user isolation.
* D only capturesmanagement events, not object-level data access.
* E is legacy, inefficient, and not structured for per-user auditing.
References:#https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_variables.html#https://docs.aws.amazon.com/awscloudtrail/latest/userguide/logging-data-events.html

**NEW QUESTION: 193**
A company is preparing to deploy an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for a workload. The company expects the cluster to support an unpredictable number of stateless pods. Many of the pods will be created during a short time period as the workload automatically scales the number of replicas that the workload uses.
Which solution will MAXIMIZE node resilience?
**A.** Use a separate launch template to deploy the EKS control plane into a second cluster that is separate from the workload node groups.
**B.** Update the workload node groups. Use a smaller number of node groups and larger instances in the node groups.
**C.** Configure the Kubernetes Cluster Autoscaler to ensure that the compute capacity of the workload node groups stays under provisioned.
**D.** Configure the workload to use topology spread constraints that are based on Availability Zone.
**Answer: D (LEAVE A REPLY)**
Configuring the workload to use topology spread constraints that are based on Availability Zone will maximize the node resilience of the workload node groups. This will ensure that the pods are evenly distributed across different Availability Zones, reducing the impact of failures or disruptions in one Availability Zone2. This will also improve the availability and scalability of the workload node groups, as they can leverage the low-latency, high-throughput, and highly redundant networking between Availability Zones1.

**NEW QUESTION: 194**
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 SOS).
**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) Me 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.
**Answer: D (LEAVE A REPLY)**
Deploying the application on Amazon EKS with managed node groups simplifies the operational overhead of managing the Kubernetes cluster. Running the web servers and application as Kubernetes deployments ensures that the desired number of pods are always running and can scale up or down as needed. Storing the frontend web server session data in an Amazon DynamoDB table provides a fast, scalable, and durable storage option that can be accessed across multiple Availability Zones. Creating an Amazon EFS volume that all applications will mount at the time of deployment allows the application to share data that is frequently accessed between the web and application tiers. References:
* https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html
* https://docs.aws.amazon.com/eks/latest/userguide/deployments.html
* https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html
* https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html

**NEW QUESTION: 195**
A company has a website that runs on four Amazon EC2 instances that are behind an Application Load Balancer (ALB). When the ALB detects that an EC2 instance is no longer available, an Amazon CloudWatch alarm enters the ALARM state. A member of the company's operations team then manually adds a new EC2 instance behind the ALB.
A solutions architect needs to design a highly available solution that automatically handles the replacement of EC2 instances. The company needs to minimize downtime during the switch to the new solution.
Which set of steps should the solutions architect take to meet these requirements?
**A.** Delete the existing ALB. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Attach the existing EC2 instances to the Auto Scaling group.

**B.** Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Attach the existing EC2 instances to the Auto Scaling group.
**C.** Delete the existing ALB and the EC2 instances. Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Create a new ALB. Attach the Auto Scaling group to the new ALB. Wait for the Auto Scaling group to launch the minimum number of EC2 instances.
**D.** Create an Auto Scaling group that is configured to handle the web application traffic. Attach a new launch template to the Auto Scaling group. Attach the Auto Scaling group to the existing ALB. Wait for the existing ALB to register the existing EC2 instances with the Auto Scaling group.
**Answer: B (LEAVE A REPLY)**
The Auto Scaling group can automatically launch and terminate EC2 instances based on the demand and health of the web application. The launch template can specify the configuration of the EC2 instances, such as the AMI, instance type, security group, and user data. The existing ALB can distribute the traffic to the EC2 instances in the Auto Scaling group. The existing EC2 instances can be attached to the Auto Scaling group without deleting them or the ALB. This option minimizes downtime and preserves the current setup of the web application. References: [What is Amazon EC2 Auto Scaling?], [Launch templates], [Attach a load balancer to your Auto Scaling group], [Attach EC2 instances to your Auto Scaling group]

**NEW QUESTION: 196**
A company needs to implement a disaster recovery (DR) plan for a web application. The application runs in a single AWS Region.
The application uses microservices that run in containers. The containers are hosted on AWS Fargate in Amazon Elastic Container Service (Amazon ECS). The application has an Amazon RDS for MYSQL DB instance as its data layer and uses Amazon Route 53 for DNS resolution. An Amazon CloudWatch alarm invokes an Amazon EventBridge rule if the application experiences a failure.
A solutions architect must design a DR solution to provide application recovery to a separate Region. The solution must minimize the time that is necessary to recover from a failure.
Which solution will meet these requirements?
**A.** Set up a second ECS cluster and ECS service on Fargate in the separate Region. Create an AWS Lambda function to perform the following actions: take asnapshot of the ROS DB instance. copy the snapshot to the separate Region. create a new RDS DB instance frorn the snapshot, and update Route 53 toroute traffic to the second ECS cluster. Update the EventBridge rule to add a target that will invoke the Lambda function.
**B.** Create an AWS Lambda function that creates a second ECS cluster and ECS service in the separate Region. Configure the Lambda function to perform thefollowing actions: take a snapshot of thQRDS DB instance, copy the snapshot to the separate Region. create a new RDS DB instance from the snapshot.and update Route 53 to route traffic to the second ECS cluster. Update the EventBridge rule to add a target that will invoke the Lambda function.

**C.** Set up a second ECS cluster and ECS service on Fargate in the separate Region. Create a cross-Region read replica of the RDS DB instance in theseparate Region. Create an AWS Lambda function to prornote the read replica to the primary database. Configure the Lambda function to update Route 53to route traffic to the second ECS cluster. Update the EventBridge rule to add a target that will invoke the Lambda function.
**D.** Set up a second ECS cluster and ECS service on Fargate in the separate Region. Take a snapshot of the ROS DB instance. Convert the snapshot to anAmazon DynamoDB global table. Create an AWS Lambda function to update Route 53 to route traffic to the second ECS cluster Update the EventBridgerule to add a target that will invoke the Lambda function.
**Answer: (SHOW ANSWER)**
This option uses a cross-Region read replica of the RDS DB instance to provide a standby database in the separate Region. A cross-Region read replica is a copy of the primary database that is updated asynchronously using the native replication features of the database engine. It provides enhanced availability, scalability, and performance for read-heavy workloads. It also enables fast recovery from a regional outage by promoting the read replica to a standalone database. To use a cross-Region read replica, the companyneeds to set up a second ECS cluster and ECS service on Fargate in the separate Region. The company also needs to create an AWS Lambda function to promote the read replica to the primary database and update Route 53 to route traffic to the second ECS cluster. The company can then update the EventBridge rule to add a target that will invoke the Lambda function in case of a failure.



**NEW QUESTION: 197**
A company has an application that runs on Amazon EC2 instances in an Amazon EC2 Auto Scaling group.
The company uses AWS CodePipeline to deploy the application. The instances that run in the Auto Scaling group are constantly changing because of scaling events.
When the company deploys new application code versions, the company installs the AWS CodeDeploy agent on any new target EC2 instances and associates the instances with the CodeDeploy deployment group. The application is set to go live within the next 24 hours.
What should a solutions architect recommend to automate the application deployment process with the LEAST amount of operational overhead?
**A.** Configure Amazon EventBridge to invoke an AWS Lambda function when a new EC2 instance is launched into the Auto Scaling group. Code the Lambda function to associate the EC2 instances with the CodeDeploy deployment group.

**B.** Write a script to suspend Amazon EC2 Auto Scaling operations before the deployment of new code When the deployment is complete, create a new AMI and configure the Auto Scaling group's launch template to use the new AMI for new launches. Resume Amazon EC2 Auto Scaling operations.
**C.** Create a new AWS CodeBuild project that creates a new AMI that contains the new code Configure CodeBuild to update the Auto Scaling group's launch template to the new AMI. Run an Amazon EC2 Auto Scaling instance refresh operation.
**D.** Create a new AMI that has the CodeDeploy agent installed. Configure the Auto Scaling group's launch template to use the new AMI. Associate the CodeDeploy deployment group with the Auto Scaling group instead of the EC2 instances.
**Answer: D (LEAVE A REPLY)**
https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html

**NEW QUESTION: 198**
A company has a data lake in Amazon S3 that needs to be accessed by hundreds of applications across many AWS accounts. The company's information security policy states that the S3 bucket must not be accessed over the public internet and that each application should have the minimum permissions necessary to function.
To meet these requirements, a solutions architect plans to use an S3 access point that is restricted to specific VPCs for each application.
Which combination of steps should the solutions architect take to implement this solution? (Select TWO.)
**A.** Create an S3 access point for each application in the AWS account that owns the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.
**B.** Create an interface endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Create a VPC gateway attachment for the S3 endpoint.
**C.** Create a gateway endpoint for Amazon S3 in each application's VPC. Configure the endpoint policy to allow access to an S3 access point. Specify the route table that is used to access the access point.
**D.** Create an S3 access point for each application in each AWS account and attach the access points to the S3 bucket. Configure each access point to be accessible only from the application's VPC. Update the bucket policy to require access from an access point.
**E.** Create a gateway endpoint for Amazon S3 in the data lake's VPC. Attach an endpoint policy to allow access to the S3 bucket. Specify the route table that is used to access the bucket.
**Answer: A,C (LEAVE A REPLY)**
https://joe.blog.freemansoft.com/2020/04/protect-data-in-cloud-with-s3-access.html

**NEW QUESTION: 199**
A telecommunications company is running an application on AWS. The company has set up anAWS 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.
**Answer: B (LEAVE A REPLY)**
The company should create a Network Load Balancer (NLB) and associate it with one static IP address in multiple Availability Zones. The company should also create an ALB-type target group for the NLB and add the existing ALB. The company should add the NLB IP addresses to the firewall appliance and update the clients to connect to the NLB. This solution will allow traffic flow to AWS from the on-premises network by using static IP addresses that can be added to the firewall appliance's allow list. The NLB will forward requests to the ALB, which will use path-based routing to forward requests to the target groups.

**NEW QUESTION: 200**
A company is deploying a distributed in-memory database on a fleet of Amazon EC2 instances. The fleet consists of a primary node and eight worker nodes. The primary node is responsible for monitoring cluster health, accepting user requests, distributing user requests to worker nodes, and sending an aggregate response back to a client. Worker nodes communicate with each other to replicate data partitions.
The company requires the lowest possible networking latency to achieve maximum performance.
Which solution will meet these requirements?
**A.** Launch memory optimized EC2 instances in a cluster placement group
**B.** Launch memory optimized EC2 instances in a partition placement group.
**C.** Launch compute optimized EC2 instances in a partition placement group.
**D.** Launch compute optimized EC2 instances in a spread placement group.
**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 201**

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:


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 Full AWS Access SCP from the developer account's OU

**C.** Modify the Full AWS Access SCP to explicitly deny all services

**D.** Add an explicit deny statement using a wildcard to the end of the SCP

**Answer: B (LEAVE A REPLY)**

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

**NEW QUESTION: 202**

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.
**Answer: B (LEAVE A REPLY)**
https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/
https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/

**NEW QUESTION: 203**
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 1AM 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.
**Answer: D (LEAVE A REPLY)**
Cross account role should be created in destination(member) account. The role has trust entity to master account.

**NEW QUESTION: 204**
To abide by industry regulations, a solutions architect must design a solution that will store a company's critical data in multiple public AWS Regions, including in the United States, where the company's headquarters is located The solutions architect is required to provide access to the data stored in AWS to the company's global WAN network The security team mandates that no traffic accessing this data should traverse the public internet How should the solutions architect design a highly available solution that meets the requirements and is cost- effective'?

**A.** Establish AWS Direct Connect connections from the company headquarters to all AWS Regions in use the company WAN to send traffic over to the headquarters and then to the respective DX connection to access the data
**B.** Establish two AWS Direct Connect connections from the company headquarters to an AWS Region Use the company WAN to send traffic over a DX connection Use inter-region VPC peering to access the data in other AWS Regions
**C.** Establish two AWS Direct Connect connections from the company headquarters to an AWS Region Use the company WAN to send traffic over a DX connection Use an AWS transit VPC solution to access data in other AWS Regions
**D.** Establish two AWS Direct Connect connections from the company headquarters to an AWS Region Use the company WAN to send traffic over a DX connection Use Direct Connect Gateway to access data in other AWS Regions.
**Answer: (SHOW ANSWER)**
* Establish AWS Direct Connect Connections:
* Step 1: Set up two AWS Direct Connect (DX) connections from the company headquarters to a chosen AWS Region. This provides a redundant and high-availability setup to ensure continuous connectivity.
* Step 2: Ensure that these DX connections terminate in a specific Direct Connect location associated with the chosen AWS Region.
* Use Company WAN:
* Step 1: Configure the company's global WAN to route traffic through the established Direct Connect connections.
* Step 2: This setup ensures that all traffic between the company's headquarters and AWS does not traverse the public internet, maintaining compliance with security requirements.
* Set Up Direct Connect Gateway:
* Step 1: Create a Direct Connect Gateway in the AWS Management Console. This gateway allows you to connect your Direct Connect connections to multiple VPCs across different AWS Regions.
* Step 2: Associate the Direct Connect Gateway with the VPCs in the various Regions where your critical data is stored. This enables access to data in multiple Regions through a single Direct Connect connection.
By using Direct Connect and Direct Connect Gateway, the company can achieve secure, reliable, and cost- effective access to data stored across multiple AWS Regions without using the public internet, ensuring compliance with industry regulations.
References
* AWS Direct Connect Documentation
* Building a Scalable and Secure Multi-VPC AWS Network Infrastructure(AWS Documentation)(AWS Documentation).

**NEW QUESTION: 205**
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 lo replicate data between the on-premises Windows file server and Amazon Elastic File System (Amazon EFS),
**Answer: B (LEAVE A REPLY)**
https://aws.amazon.com/storagegateway/file/
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html
https://docs.aws.amazon.com/systems-manager/latest/userguide/prereqs-operating-systems.html#prereqs-os- windows-server

**NEW QUESTION: 206**
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#¢#' 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 VPC. Attach 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 image on the EFS volume.
**Answer: (SHOW ANSWER)**
Create Amazon S3 gateway endpoint in the VPC and add a VPC endpoint policy. This VPC endpoint policy will have a statement that allows S3 access only via access points owned by the organization.

**NEW QUESTION: 207**

A company deploys workloads in multiple AWS accounts. Each account has a VPC with VPC flow logs published in text log format to a centralized Amazon S3 bucket. Each log file is compressed with gzjp compression. The company must retain the log files indefinitely.
A security engineer occasionally analyzes the togs by using Amazon Athena to query the VPC flow logs. The query performance is degrading over time as the number of ingested togs is growing. A solutions architect:
must improve the performance of the tog analysis and reduce the storage space that the VPC flow logs use.
Which solution will meet these requirements with the LARGEST performance improvement?
**A.** Create an AWS Lambda function to decompress the gzip flies and to compress the tiles with bzip2 compression. Subscribe the Lambda function to an s3: ObiectCrealed;Put S3 event notification for the S3 bucket.
**B.** Enable S3 Transfer Acceleration for the S3 bucket. Create an S3 Lifecycle configuration to move files to the S3 Intelligent-Tiering storage class as soon as the ties are uploaded
**C.** Update the VPC flow log configuration to store the files in Apache Parquet format. Specify Hourly partitions for the log files.
**D.** Create a new Athena workgroup without data usage control limits. Use Athena engine version 2.
**Answer: (SHOW ANSWER)**
Converting VPC flow logs to store in Apache Parquet format and specifying hourly partitions significantly improves query performance and reduces storage space usage. Apache Parquet is a columnar storage file format optimized for analytical queries, allowing Athena to scan less data and improve query performance.
Partitioning logs by hour further enhances query efficiency by limiting the amount of data scanned during queries, addressing the issue of degrading performance over time due to the growing volume of ingested logs.
AWS Documentation on VPC Flow Logs and Amazon Athena provides insights into configuring VPC flow logs in Apache Parquet format and using Athena for querying log data. This approach is recommended for efficient log analysis and storage optimization.

**NEW QUESTION: 208**
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 Replicationbetween 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.
**Answer: B (LEAVE A REPLY)**
an AWS Lambda function in the backup region to promote the read replica and modify the Auto Scaling group values, and then configuring Route 53 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. Finally, the application's Route 53 record should be updated with a failover policy that routes traffic to the ALB in the backup region when a health check failure occurs. This approach provides automatic failover to the backup region when a health check failure occurs, reducing the RTO to less than 15 minutes. Additionally, this approach is cost-effective as it does not require an active-active strategy.

**NEW QUESTION: 209**
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
**Answer: A,D (LEAVE A REPLY)**
The company should use Amazon Aurora global database and Amazon DynamoDB global table to deploy the data tier components across two Regions. Amazon Aurora globaldatabase is a feature that allows a single Aurora database to span multiple AWS Regions, enabling low-latency global reads and fast recovery from Region-wide outages1. Amazon DynamoDB global table is a feature that allows a single DynamoDB table to span multiple AWS Regions, enabling low-latency global reads and writes and fast recovery from Region- wide outages2.
References:
https://aws.amazon.com/rds/aurora/global-database/
https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/globaltables_HowItWorks.html
https://aws.amazon.com/route53/application-recovery-controller/

**NEW QUESTION: 210**
A company's solutions architect is evaluating an AWS workload that was deployed several years ago. The application tier is stateless and runs on a single large Amazon EC2 instance that was launched from an AMI.
The application stores data in a MySOL database that runs on a single EC2 instance.
The CPU utilization on the application server EC2 instance often reaches 100% and causes the application to stop responding. The company manually installs patches on the instances. Patching has caused downtime in the past. The company needs to make the application highly available.
Which solution will meet these requirements with the LEAST development time?
**A.** Move the application tier to AWS Lambda functions in the existing VPC. Create an Application Load Balancer to distribute traffic across the Lambda functbns. Use Amazon GuardDuty to scan the Lambda functions. Migrate the database to Amazon DocumentDB (with MongoDB compatibility).
**B.** Change the EC2 instance type to a smaller Graviton powered instance type. use the existing AMI to create a launch template for an Auto Scaling group. Create an Application Load Balancer to distribute traffic across the instances in the Auto Scaling group. Set the Auto Scaling group to scale based on CPU utilization. Migrate the database to Amazon DynamoDB.
**C.** Move the application tier to containers by using Docker. Run the containers on Amazon Elastic Container Service (Amazon ECS) with EC2 instances. Create an Application Load Balancer to distribute traffic across the ECS cluster Configure the ECS cluster to scale based on CPU utilization.
Migrate the database to Amazon Neptune.
**D.** Create a new AMI that is configured with AWS Systems Manager Agent (SSM Agent). Use the new AMI to create a launch template for an Auto Scaling group. Use smaller instances in the Auto Scaling group. Create an Application Load Balancer to distribute traffic across the instances in the Auto Scaling group. Set the Auto Scaling group to scale based on CPU utilization. Migrate the database to Amazon Aurora MySQL.

**Answer: D (LEAVE A REPLY)**

This solution will meet the requirements of making the application highly available with the least development time. Creating a new AMI that is configured with SSM Agent will enable the company to use AWS Systems Manager to manage and patch the EC2 instances automatically, reducing downtime and human errors. Using a launchtemplate for an Auto Scaling group will allow the company to launch multiple instances of the same configuration and scale them up or down based on demand. Using smaller instances in the Auto Scaling group will reduce the cost and improve the performance of the application tier. Creating an Application Load Balancer to distribute traffic across the instances in the Auto Scaling group will increase the availability and fault tolerance of the application tier. Migrating the database to Amazon Aurora MySQL will provide a fully managed, compatible, and scalable relational database service that can handle high throughput and concurrent connections.