Amazon.AWS-Solutions-Architect-Professional.v2024-10-15.q255

Exam Code:	AWS-Solutions-Architect-Professional
Exam Name:	AWS Certified Solutions Architect - Professional
Certification Provider:	Amazon
Free Question Number:	255
Version:	v2024-10-15
# of views:	2289
# of Questions views:	6273
https://www.dumpsfiles.com/files/Amazon/AWS-Solutions-Architect-Professional/Amazon.AWS-Solutions-Architect-Professional.v2024-10-15.q255	

NEW QUESTION: 1
You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location.
What options could you select to migrate the application to AWS? (Choose 2)
A. Create Dockerfile for the application. Create an AWS OpsWorks stack consisting of a custom layer. Create custom recipes to install Docker and to deploy your Docker container using the Dockerfile. Create customer recipes to install and configure the application to publish the logs to Amazon CloudWatch Logs.
B. Create a Dockerfile for the application. Create an AWS Elastic Beanstalk application using the Docker platform and the Dockerfile. Enable logging the Docker configuration to automatically publish the application logs. Enable log file rotation to Amazon S3.
C. Create an AWS Elastic Beanstalk application using the custom web server platform. Specify the web server executable and the application project and source files. Enable log file rotation to Amazon Simple Storage Service (S3).
D. Use VM import/Export to import a virtual machine image of the server into AWS as an AMI. Create an Amazon Elastic Compute Cloud (EC2) instance from AMI, and install and configure the Amazon CloudWatch Logs agent. Create a new AMI from the instance. Create an AWS Elastic Beanstalk application using the AMI platform and the new AMI.
E. Create Dockerfile for the application. Create an AWS OpsWorks stack consisting of a Docker layer that uses the Dockerfile. Create custom recipes to install and configure Amazon Kineses to publish the logs into Amazon CloudWatch.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 2
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: B (LEAVE A REPLY)
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_

NEW QUESTION: 3
What is the average queue length recommended by AWS to achieve a lower latency for the 200 PIOPS EBS volume?
A. 5
B. 1
C. 2
D. 4
Answer: B (LEAVE A REPLY)
The queue length is the number of pending I/O requests for a device. The optimal average queue length will vary for every customer workload, and this value depends on a particular application's sensitivity to IOPS and latency. If the workload is not delivering enough I/O requests to maintain the optimal average queue length, then the EBS volume might not consistently deliver the IOPS that have been provisioned. However, if the workload maintains an average queue length that is higher than the optimal value, then the per-request I/O latency will increase; in this case, the user should provision more IOPS for his volume. AWS recommends that the user should target an optimal average queue length of 1 for every 200 provisioned IOPS and tune that value based on his application requirements.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-workload-demand.html

NEW QUESTION: 4
A company is running an application in a single VPC on an Amazon EC2 instance with Amazon RDS as the datastore. The application does not support encryption in transit Security guidelines do not allow SSH access to any resource within the VPC.

The Application has issues throughout the day which causes outages in the production environment. The issues are not present in nonproduction environments Application logs have been given to a vendor to troubleshoot the application. The vendor also requires IP packets for its analysis.
Which solution allows for the IP packets to be extracted for troubleshooting?
A. Create a VPC traffic mirror source on the application instance's elastic network interface with a filter that captures all traffic. Configure the traffic mirror target to use an Amazon S3 bucket Start the traffic mirror session and download the packet capture from Amazon S3. Provide the packet capture to the vendor.
B. Enable VPC Flow Logs on the application instance's elastic network interface and send them to Amazon CloudWatch Logs Download the CloudWatch logs and provide them to me vendor
C. Create a VPC traffic mirror source on the application instance's elastic network interface with a filter that captures all traffic Launch a new EC2 instance and configure the traffic minor target to use the elastic network interface of the new EC2 instance. Start the traffic mirror session and download the packet capture from the new EC2 instance using AWS Systems Manager Provide the packet capture to the vendor.
D. Enable VPC Flow Logs on the VPC to capture traffic flows on from the application instance and the RDS instance and send them to Amazon CloudWatch Logs Download the CloudWatch logs and provide them to the vendor
Answer: D (LEAVE A REPLY)

NEW QUESTION: 5
A company is running several workloads in a single AWS account. A new company policy states that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.
What should the solutions architect do to create the solution?
A. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket.
Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormation.
Use AWS CloudFormation templates to provision resources.
B. Update the IAM policy for the engineers' IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormation. Use AWS CloudFormation templates to create stacks with approved resources.
C. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role. Assign the IAM service role to AWS CloudFormation during stack creation.
D. Provision resources in AWS CloudFormation stacks. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.
Answer: B (LEAVE A REPLY)
Explanation
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/security-best-practices.html#use-iam-to-co
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

NEW QUESTION: 6

A company wants to host a new global website that consists of static content. A solutions architect is working on a solution that uses Amazon CloudFront with an origin access identity <OAI) to access website content that is stored in a private Amazon S3 bucket.
During testing, the solutions architect receives 404 errors from the S3 bucket. Error messages appear only for attempts to access paths that end with a forward slash. such as example.com/path/. These requests should return the existing S3 object path/index.html. Any potential solution must not prevent CloudFront from caching the content.
What should the solutions architect do to resolve this problem?
A. Change the CloudFront origin to an Amazon API Gateway endpoint. Rewrite the S3 request URL in an AWS service integration.
B. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by an origin request event to rewrite the S3 request URL.
C. Change the CloudFront origin to an Amazon API Gateway proxy endpoint. Rewrite the S3 request URL by using an AWS Lambda function.
D. Change the CloudFront configuration to use an AWS Lambda@Edge function that is invoked by a viewer request event to rewrite the S3 request URL.
Answer: (SHOW ANSWER)

NEW QUESTION: 7
Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24.
While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance.
Which is the most likely reason for this issue?
A. Private address IP 10.201.31.6 is currently assigned to another interface.
B. Private IP address 10.201.31.6 is reserved by Amazon for IP networking purposes.
C. Private IP address 10.201.31.6 is blocked via ACLs in Amazon infrastructure as a part of platform security.
D. Private IP address 10.201.31.6 is not part of the associated subnet's IP address range.
Answer: A (LEAVE A REPLY)
In Amazon VPC, you can assign any Private IP address to your instance as long as it is:
Part of the associated subnet's IP address range
Not reserved by Amazon for IP networking purposes
Not currently assigned to another interface
http://aws.amazon.com/vpc/faqs/

NEW QUESTION: 8
A financial services company logs personally identifiable information 10 its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The security team has mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.
Which steps should the solutions architect take to meet these requirements?

A. Create an AWS CloudHSM cluster. Create a new CMK in AWS KMS using AWS_CloudHSM as the source (or the key material and an origin of AWS_CLOUDHSM. Enable automatic key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the togging bucket thai disallows uploads of unencrypted data and requires that the encryption source be AWS KMS.
B. Provision an AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPCs. Configure an AWS bucket policy on the logging bucket that requires all objects to be encrypted. Configure the logging application to query the on-premises HSMs from the AWS environment for the encryption key material, and create a unique CMK for each logging event.
C. Create a CMK in AWS KMS with no key material and an origin of EXTERNAL. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
D. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS_KMS.
Disable this CMK. and overwrite the key material with the key material from the on-premises HSM using the public key and import token provided by AWS. Re-enable the CMK. Enable automatic key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
Answer: (SHOW ANSWER)
Explanation
https://aws.amazon.com/blogs/security/how-to-byok-bring-your-own-key-to-aws-kms-for-less-than-15-00-a-year
https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-create-cmk.html

NEW QUESTION: 9
You are playing around with setting up stacks using JSON templates in CloudFormation to try and understand them a little better. You have set up about 5 or 6 but now start to wonder if you are being charged for these stacks. What is AWS's billing policy regarding stack resources?
A. You are not charged for the stack resources if they are not taking any traffic.
B. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 30 minutes)
C. You are charged for the stack resources for the time they were operating (but not if you deleted the stack within 60 minutes)
D. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away)
Answer: D (LEAVE A REPLY)
A stack is a collection of AWS resources that you can manage as a single unit. In other words, you can create, update, or delete a collection of resources by creating, updating, or deleting stacks. All the resources in a stack are defined by the stack's AWS CloudFormation template. A stack, for instance, can include all the resources required to run a web application, such as a web server, a database, and networking rules. If you no longer require that web application, you can simply delete the stack, and all of its related resources are deleted. You are charged for the stack resources for the time they were operating (even if you deleted the stack right away).
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/stacks.html

NEW QUESTION: 10
Which of the following statements is correct about AWS Direct Connect?
A. Connections to AWS Direct Connect require double clad fiber for 1 gigabit Ethernet with Auto Negotiation enabled for the port.
B. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with.
C. AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 50 gigabit Ethernet cable.
D. To use AWS Direct Connect, your network must be colocated with a new AWS Direct Connect location.
Answer: (SHOW ANSWER)
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. To use AWS Direct Connect, your network is colocated with an existing AWS Direct Connect location. Connections to AWS Direct Connect require single mode fiber,
1000BASE-LX (1310nm) for 1 gigabit Ethernet, or 10GBASE-LR (1310nm) for 10 gigabit Ethernet. Auto Negotiation for the port must be disabled.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

NEW QUESTION: 11
A software as a service (SaaS) based company provides a case management solution to customers A3 part of the solution. The company uses a standalone Simple Mail Transfer Protocol (SMTP) server to send email messages from an application. The application also stores an email template for acknowledgement email messages that populate customer data before the application sends the email message to the customer.
The company plans to migrate this messaging functionality to the AWS Cloud and needs to minimize operational overhead.
Which solution will meet these requirements MOST cost-effectively?
A. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.
B. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template in an Amazon S3 bucket. Create an AWS Lambda function to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. Use an SDK in the Lambda function to send the email message.
C. Set up an SMTP server on Amazon EC2 instances by using an AMI from the AWS Marketplace. Store the email template in Amazon Simple Email Service (Amazon SES) with parameters for the customer data. Create an AWS Lambda function to call the SES template and to pass customer data to replace the parameters. Use the AWS Marketplace SMTP server to send the email message.
D. Set up Amazon Simple Email Service (Amazon SES) to send email messages. Store the email template on Amazon SES with parameters for the customer data. Create an AWS Lambda function to call the

SendTemplatedEmail API operation and to pass customer data to replace the parameters and the email destination.

Answer: B (LEAVE A REPLY)

Explanation

This solution would meet the company's requirements most cost-effectively because it uses Amazon SES to send email messages, which is a fully managed service that eliminates the need to set up and maintain an SMTP server. The email template is stored in an Amazon S3 bucket, which is a cost-effective object storage service that is highly durable and available. An AWS Lambda function can be used to retrieve the template from the S3 bucket and to merge the customer data from the application with the template. The Lambda function can also use an SDK to send the email message. This approach eliminates the need for EC2 instances and it's more cost-effective than other options.

You can refer to the Amazon SES documentation for more information on how to use this service:

https://aws.amazon.com/ses/

Amazon SES is a cost-effective and scalable email service that allows you to send transactional and marketing emails. Additionally, you can use SES to send templated email, which can be stored in the S3 bucket and retrieve by Lambda function, which will reduce the operational overhead. You can also use SNS for triggering the Lambda function when a new email need to be sent.

NEW QUESTION: 12

One of the components that is part of ec2-net-utils used with ENI's is ec2ifscan. Which of the following is not correct about ec2-net-utils?

A. ec2-net-utils generates an interface configuration file suitable for use with DHCP.

B. ec2-net-utils extends the functionality of the standard ifup.

C. ec2-net-utils detaches a primary network interface from an instance.

D. ec2-net-utils identifies network interfaces when they are attached, detached, or reattached to a running instance.

Answer: C (LEAVE A REPLY)

Each instance in a VPC has a default elastic network interface (the primary network interface) that is assigned a private IP address from the IP address range of your VPC. You cannot detach a primary network interface from an instance. You can create and attach additional elastic network interfaces. Amazon Linux AMIs may contain additional scripts installed by AWS, known as ec2-net-utils. One of the components that is part of ec2-net-utils used with ENI's is ec2ifscan.

Its function is to check for network interfaces that have not been configured and configure them.

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

NEW QUESTION: 13

An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.

If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP

of the instance, how can they achieve this?
A. The organization should ensure that the IP which is required to receive the ELB traffic is attached to a primary network interface.
B. It is not possible to attach an instance with two ENIs with ELB as it will give an IP conflict error.
C. The organization should ensure that the IP which is required to receive the ELB traffic is attached to an additional ENI.
D. It is not possible to send data to a particular IP as ELB will send to any one EIP.
Answer: A (LEAVE A REPLY)
Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet.
When the user registers a multi-homed instance (an instance that has an Elastic Network Interface (ENI) attached) with a load balancer, the load balancer will route the traffic to the IP address of the primary network interface (eth0).
Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/gs-ec2VPC.html

NEW QUESTION: 14
A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.
What should a solutions architect do to meet these requirements?
A. Create a new developer account. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organizations. Enforce a tagging policy that denotes Region affinity.
B. Create an IAM policy than allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.
C. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.
D. Create an SCP that denies the launch of all EC2 instances except I3.small EC2 instances in us-east-2. Attach the SCP to the project's account.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 15
A company hosts its website on Amazon S3. The website serves petabytes of outbound traffic monthly, which accounts for most of the company's AWS costs. What should a solutions architect do to reduce costs?
A. Rearchitect the website to run on a combination of Amazon API Gateway and AWS Lambda.
B. Configure Amazon CloudFront with the existing website as the origin.
C. Use AWS Global Accelerator and specify the existing website as the endpoint.
D. Move the website to Amazon EC2 with Amazon EBS volumes for storage.

Answer: (SHOW ANSWER)

NEW QUESTION: 16
A company's CISO has asked a Solutions Architect to re-engineer the company's current CI/CD practices to make sure patch deployments to its applications can happen as quickly as possible with minimal downtime if vulnerabilities are discovered. The company must also be able to quickly roll back a change in case of errors. The web application is deployed in a fleet of Amazon EC2 instances behind an Application Load Balancer. The company is currently using GitHub to host the application source code, and has configured an AWS CodeBuild project to build the application. The company also intends to use AWS CodePipeline to trigger builds from GitHub commits using the existing CodeBuild project.
What CI/CD configuration meets all of the requirements?
A. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for in-place deployment. Monitor the newly deployed code, and, if there are any issues, push another code update.
B. Configure the CodePipeline with a deploy stage using AWS OpsWorks and in-place deployments. Monitor the newly deployed code, and, if there are any issues, push another code update.
C. Configure CodePipeline with a deploy stage using AWS CloudFormation to create a pipeline for test and production stacks. Monitor the newly deployed code, and, if there are any issues, push another code update.
D. Configure CodePipeline with a deploy stage using AWS CodeDeploy configured for blue/green deployments. Monitor the newly deployed code, and, if there are any issues, trigger a manual rollback using CodeDeploy.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 17
A company is backing up on-premises databases to local file server shares using the SMB protocol. The company requires immediate access to 1 week of backup files to meet recovery objectives. Recovery after a week is less likely to occur, and the company can tolerate a delay in accessing those older backup files.
What should a solutions architect do to meet these requirements with the LEAST operational effort?
A. Deploy Amazon Elastic File System (Amazon EFS) to create a file system with exposed NFS shares with sufficient storage to hold all the desired backups.
B. Deploy Amazon FSx for Windows File Server to create a file system with exposed file shares with sufficient storage to hold all the desired backups.

C. Deploy an AWS Storage Gateway file gateway with sufficient storage to hold 1 week of backups. Point the backups to SMB shares from the file gateway.
D. Continue to back up to the existing file shares. Deploy AWS Database Migration Service (AWS DMS) and define a copy task to copy backup files older than 1 week to Amazon S3, and delete the backup files from the local file store.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 18
A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled.
Which solution will meet these requirements?
A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key. Use the AWS CLI to re-upload all objects in the S3 bucket. Set an S3 bucket policy to deny unencrypted PutObject requests.
B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests.
Use the AWS CLI to re-upload all objects in the S3 bucket.
C. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.
D. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key.
Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket. Use the AWS CLI to re-upload all objects in the S3 bucket.
Answer: B (LEAVE A REPLY)
Explanation
This solution meets the requirement that all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages by using Server-Side Encryption with AWS KMS Managed Keys (SSE-KMS). It allows the company's security team to manage, rotate and track the encryption keys and also encrypts all objects in the S3 bucket at rest. To implement this solution, you need to change the default encryption in the S3 bucket properties to SSE-KMS, by doing this all the current and future objects in the bucket will be encrypted by keys that the company's security team manages. Then set an S3 bucket policy to deny unencrypted PutObject requests. Finally, use the AWS CLI to re-upload all objects in the S3 bucket.
It's important to note that since the bucket does not have versioning enabled, re-uploading all objects will overwrite the existing objects in the bucket.
References:
* https://aws.amazon.com/s3/features/server-side-encryption/
* https://aws.amazon.com/kms/
* https://aws.amazon.com/s3/features/server-side-encryption/
* https://aws.amazon.com/s3/faqs/

* https://docs.aws.amazon.com/AmazonS3/latest/user-guide/encryption-using-server-side-encryption.html

NEW QUESTION: 19
An organization has 4 people in the IT operations team who are responsible to manage the AWS infrastructure. The organization wants to setup that each user will have access to launch and manage an instance in a zone which the other user cannot modify. Which of the below mentioned options is the best solution to set this up?
A. Create four AWS accounts and give each user access to a separate account.
B. Create an IAM user and allow them permission to launch an instance of a different sizes only.
C. Create four IAM users and four VPCs and allow each IAM user to have access to separate VPCs.
D. Create a VPC with four subnets and allow access to each subnet for the individual IAM user.
Answer: D (LEAVE A REPLY)
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. The user can create subnets as per the requirement within a VPC. The VPC also work with IAM and the organization can create IAM users who have access to various VPC services. The organization can setup access for the IAM user who can modify the security groups of the VPC. The sample policy is given below:
{
"Version": "2012-10-17",
"Statement":
[{ "Effect": "Allow",
"Action": "ec2:RunInstances",
"Resource":
["arn:aws:ec2:region::image/ami-*",
"arn:aws:ec2:region:account:subnet/subnet-1a2b3c4d",
"arn:aws:ec2:region:account:network-interface/*",
"arn:aws:ec2:region:account:volume/*",
"arn:aws:ec2:region:account:key-pair/*",
"arn:aws:ec2:region:account:security-group/sg-123abc123"]
}]
}
With this policy the user can create four subnets in separate zones and provide IAM user access to each subnet.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IAM.html

NEW QUESTION: 20
An organization is planning to use NoSQL DB for its scalable data needs. The organization wants to host an application securely in AWS VPC.
What action can be recommended to the organization?
A. The organization should setup their own NoSQL cluster on the AWS instance and configure route tables and subnets.

B. The organization should only use a DynamoDB because by default it is always a part of the default subnet provided by AWS.
C. The organization should use a DynamoDB while creating a table within the public subnet.
D. The organization should use a DynamoDB while creating a table within a private subnet.
Answer: (SHOW ANSWER)
Explanation
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Currently VPC does not support DynamoDB. Thus, if the user wants to implement VPC, he has to setup his own NoSQL DB within the VPC.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Introduction.html

NEW QUESTION: 21
An IT company owns a web product in AWS that provides discount restaurant information to customers. It has used one S3 Bucket (my_bucket) to store restaurant data such as pictures, menus, etc. The product is deployed in VPC subnets. The company's Cloud Architect decides to configure a VPC endpoint for this S3 bucket so that the performance will be enhanced. To be compliance to security rules, it is required that the new VPC endpoint is only used to communicate with this specific S3 Bucket and on the other hand, the S3 bucket only allows the read/write operations coming from this VPC endpoint.
Which two options should the Cloud Architect choose to meet the security needs?
A. In the S3 bucket "my_bucket", add a S3 bucket policy in which all actions are denied if the source IP address is not equal to the EC2 public IP (use "NotIpAddress" condition).
B. Modify the security group of the EC2 instance to limit the outbound actions to the VPC Endpoint if the outgoing traffic destination is the S3 bucket "my_bucket'".
C. For the S3 bucket "my_bucket", use a S3 bucket policy that denies all actions if the source VPC Endpoint is no equal to the endpoint ID that is created.
D. Create a S3 bucket policy in the S3 bucket "my_bucket" which denies all actions unless the source IP address is equal to the EC2 public IP (use "IpAddress" condition).
E. Use a VPC Endpoint policy for Amazon S3 to restrict access to the S3 Bucket "my_bucket" so that the VPC Endpoint is only allowed to perform S3 actions on "my_bucket".
Answer: (SHOW ANSWER)

NEW QUESTION: 22
A healthcare company runs a production workload on AWS that stores highly sensitive personal information. The security team mandates that, for auditing purposes, any AWS API action using AWS account root user credentials must automatically create a high-priority ticket in the company's ticketing system. The ticketing system has a monthly 3-hour maintenance window when no tickets can be created.
To meet security requirements, the company enabled AWS CloudTrail logs and wrote a scheduled AWSLambda function that uses Amazon Athena to query API actions performed by the root user. The Lambda function submits any actions found to the ticketing system API. During a recent security audit, the security team

discovered that several tickets were not created because the ticketing system was unavailable due to planned maintenance.
Which combination of steps should a solutions architect take to ensure that the incidents are reported to the ticketing system even during planned maintenance? (Select TWO.)
A. Modify the Lambda function to be triggered by messages published to an Amazon SNS topic. Update the existing application code to retry every 5 minutes if the ticketing system's API endpoint is unavailable.
B. Create an Amazon SQS queue to which Amazon CloudWatch alarms will be published. Configure a CloudWatch alarm to publish to the SQS queue.
C. Create an Amazon SNS topic to which Amazon CloudWatch alarms will be published. Configure a CloudWatch alarm to invoke the Lambda function.
D. Create an Amazon EventBridge rule that triggers on all API events where the invoking user identity is root. Configure the EventBridge rule to write the event to an Amazon SQS queue.
E. Modify the Lambda function to be triggered when there are messages in the Amazon SQS queue and to return successfully when the ticketing system API has processed the request.
Answer: B,E (LEAVE A REPLY)

NEW QUESTION: 23
You have an application running on an EC2 instance which will allow users to download files from a private S3 bucket using a pre-signed URL. Before generating the URL, the application should verify the existence of the file in S3.
How should the application use AWS credentials to access the S3 bucket securely?
A. Use the AWS account access keys; the application retrieves the credentials from the source code of the application.
B. Create an IAM role for EC2 that allows list access to objects In the S3 bucket; launch the Instance with the role, and retrieve the role's credentials from the EC2 instance metadata.
C. Create an IAM user for the application with permissions that allow list access to the S3 bucket; the application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the Application user.
D. Create an IAM user for the application with permissions that allow list access to the S3 bucket; launch the instance as the IAM user, and retrieve the IAM user's credentials from the EC2 instance user data.
Answer: B (LEAVE A REPLY)
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html

NEW QUESTION: 24
A user is accessing an EC2 instance on the SSH port for IP 10.20.30.40/32.
Which one is a secure way to configure that the instance can be accessed only from this IP?
A. In the security group, open port 22 for IP 10.20.30.40
B. In the security group, open port 22 for IP 10.20.30.0
C. In the security group, open port 22 for IP 10.20.30.40/32
D. In the security group, open port 22 for IP 10.20.30.40/0
Answer: C (LEAVE A REPLY)

Explanation
In AWS EC2, while configuring a security group, the user needs to specify the IP address in CIDR notation. The CIDR IP range 10.20.30.40/32 says it is for a single IP 10.20.30.40. If the user specifies the IP as 10.20.30.40 only, the security group will not accept and ask for it in a CIDR format.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html

NEW QUESTION: 25
When I/O performance is more important than fault tolerance, which of the following configurations should be used?
A. SPAN 10
B. RAID 1
C. RAID 0
D. NFS 1
Answer: C (LEAVE A REPLY)
Explanation
When I/O performance is more important than fault tolerance, the RAID 0 configuration must be used; for example, as in a heavily used database (where data replication is already set up separately).
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

NEW QUESTION: 26
A company's AWS architecture currently uses access keys and secret access keys stored on each instance to access AWS services Database credentials are hard-coded on each instance SSH keys for command-line remote access are stored in a secured Amazon S3 bucket The company has asked its solutions architect to improve the security posture of the architecture without adding operational complexly.
Which combination of steps should the solutions architect take to accomplish this? (Select THREE.)
A. Use AWS Systems Manager Parameter Store to store database credentials
B. Use AWS KMS to store database credentials
C. Use a secure fleet of Amazon EC2 bastion hosts for remote access
D. Use AWS Systems Manager Session Manager for remote access
E. Use Amazon EC2 instance profiles with an IAM role
F. Use AWS Secrets Manager to store access keys and secret access keys
Answer: (SHOW ANSWER)

NEW QUESTION: 27
Which statement is NOT true about accessing remote AWS region in the US by your AWS Direct Connect which is located in the US?
A. AWS Direct Connect locations in the United States can access public resources in any US region.
B. You can use a single AWS Direct Connect connection to build multi-region services.
C. Any data transfer out of a remote region is billed at the location of your AWS Direct Connect data transfer rate.

D. To connect to a VPC in a remote region, you can use a virtual private network (VPN) connection over your public virtual interface.

Answer: C (LEAVE A REPLY)

Explanation

AWS Direct Connect locations in the United States can access public resources in any US region. You can use a single AWS Direct Connect connection to build multi-region services. To connect to a VPC in a remote region, you can use a virtual private network (VPN) connection over your public virtual interface.

To access public resources in a remote region, you must set up a public virtual interface and establish a border gateway protocol (BGP) session. Then your router learns the routes of the other AWS regions in the US. You can then also establish a VPN connection to your VPC in the remote region.

Any data transfer out of a remote region is billed at the remote region data transfer rate.

http://docs.aws.amazon.com/directconnect/latest/UserGuide/remote_regions.html

NEW QUESTION: 28

A company's security compliance requirements state that all Amazon EC2 images must be scanned for vulnerabilities and must pass a CVE assessment. A solutions architect is developing a mechanism to create security- approved AMIs that can be used by developers. Any new AMIs should go through an automated assessment process and be marked as approved before developers can use them. The approved images must be scanned every30 days to ensure compliance.

Which combination of steps should the solutions architect take to meet these requirements while following best practices? (Select TWO.)

A. Use the AWS Systems Manager EC2 agent to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned.

B. Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use a managed AWS Config rule for continuous scanning on all EC2 instances, and use AWS Systems Manager Automation documents for remediation.

C. Use AWS CloudTrail to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned.

D. Use Amazon Inspector to run the CVE assessment on the EC2 instances launched from the AMIs that need to be scanned.

E. Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days.

Answer: (SHOW ANSWER)

NEW QUESTION: 29

A solutions architect needs to improve an application that is hosted in the AWS Cloud. The application uses an Amazon Aurora MySQL DB instance that is experiencing overloaded connections. Most of the application's operations insert records into the database. The application currently stores credentials in a text-based configuration file.

The solutions architect needs to implement a solution so that the application can handle the current connection load. The solution must keep the credentials secure and must provide the ability to rotate the credentials automatically on a regular basis.
Which solution will meet these requirements?
A. Deploy an Amazon RDS Proxy layer in front of the DB instance. Store the connection credentials as a secret in AWS Secrets Manager.
B. Deploy an Amazon RDS Proxy layer in front of the DB instance. Store the connection credentials in AWS Systems Manager Parameter Store.
C. Create an Aurora Replica. Store the connection credentials as a secret in AWS Secrets Manager.
D. Create an Aurora Replica. Store the connection credentials in AWS Systems Manager Parameter Store.
Answer: A (LEAVE A REPLY)
Explanation
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/rds-proxy.html

NEW QUESTION: 30
A company wants to allow its Marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The Team Manager must have the ability to manage users and groups, but no team members should have access to services or resources not required for the SQL queries. Additionally, Administrators need to audit the queries made and receive notifications when a query violates rules defined by the Security team.
AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the Team Manager.
Which design meets these requirements?
A. Apply a service control policy (SCP) that denies to all services except IAM, Amazon DynamoDB, and AWS CloudTrail. Store customer records in DynamoDB and train users to execute queries using the AWS CLI. Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting.
B. Apply a service control policy (SCP) that allows access to IAM, Amazon RDS, and AWS CloudTrail. Load customer records in Amazon RDS MySQL and train users to execute queries using the AWS CLI. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance. use a subscription filter with AWS lambda functions to audit and alarm on queries against personal data.
C. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon Athena, Amazon S3, and AWS CloudTrail. Store customer record files in Amazon S3 and train users to execute queries using the CLI via Athena. Analyze CloudTrail events to audit and alarm on queries against personal data.
D. Apply a service control policy (SCP) that allows to IAM, Amazon Athena, Amazon S3, and AWS CloudTrail. Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and execute queries using the AWS CLI. Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 31

A customer has a website which shows all the deals available across the market. The site experiences a load of 5 large EC2 instances generally. However, a week before Thanksgiving vacation they encounter a load of almost 20 large instances. The load during that period varies over the day based on the office timings. Which of the below mentioned solutions is cost effective as well as help the website achieve better performance?

A. Setup to run 10 instances during the pre-vacation period and only scale up during the office time by launching 10 more instances using the AutoScaling schedule.

B. Keep only 10 instances running and manually launch 10 instances every day during office hours.

C. During the pre-vacation period setup 20 instances to run continuously.

D. During the pre-vacation period setup a scenario where the organization has 15 instances running and 5 instances to scale up and down using Auto Scaling based on the network I/O policy.

Answer: D (LEAVE A REPLY)

AWS provides an on demand, scalable infrastructure. AWS EC2 allows the user to launch On-Demand instances and the organization should create an AMI of the running instance. When the organization is experiencing varying loads and the time of the load is not known but it is higher than the routine traffic it is recommended that the organization launches a few instances before hand and then setups AutoScaling with policies which scale up and down as per the EC2 metrics, such as Network I/O or CPU utilization.

If the organization keeps all 10 additional instances as a part of the AutoScaling policy sometimes during a sudden higher load it may take time to launch instances and may not give an optimal performance. This is the reason it is recommended that the organization keeps an additional 5 instances running and the next 5 instances scheduled as per the AutoScaling policy for cost effectiveness.

Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION: 32

A company has an application that runs on Amazon EC2 instances. A solutions architect is designing VPC infrastructure in an AWS Region where the application needs to access an Amazon Aurora DB cluster. The EC2 instances are all associated with the same security group. The DB cluster is associated with its own security group.

The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB cluster.

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

A. Add an inbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the source over the default Aurora port.
B. Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port.
C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port.
D. Add an outbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the destination over the default Aurora port.
E. Add an outbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the destination over the ephemeral ports.
Answer: B,C (LEAVE A REPLY)
Explanation
B) Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port. This allows the instances to make outbound connections to the DB cluster on the default Aurora port. C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port. This allows connections to the DB cluster from the EC2 instances on the default Aurora port.

NEW QUESTION: 33
An ecommerce website running on AWS uses an Amazon RDS for MySQL DB instance with General Purpose SSD storage. The developers chose an appropriate instance type based on demand, and configured
100 GB of storage with a sufficient amount of free space.
The website was running smoothly for a few weeks until a marketing campaign launched On the second day of the campaign, users reported long wait times and time outs Amazon CloudWatch metrics indicated that both reads and writes to the DB instance were experiencing long response times. The CloudWatch metrics show 40% to 50% CPU and memory utilization, and sufficient free storage space is still available The application server logs show no evidence of database connectivity issues What could be the root cause of the issue with the marketing campaign?
A. It exhausted the network bandwidth available to the RDS for MySQL DB instance
B. It caused the data in the tables to change frequently, requiring indexes to be rebuilt to optimize queries
C. It exhausted the I/O credit balance due to provisioning low disk storage during the setup phase
D. It exhausted the maximum number of allowed connections to the database instance
Answer: C (LEAVE A REPLY)

NEW QUESTION: 34
Identify a correct statement about the expiration date of the "Letter of Authorization and Connecting Facility Assignment (LOA-CFA)," which lets you complete the Cross Connect step of setting up your AWS Direct Connect.
A. If the cross connect is not completed within 90 days, the authority granted by the LOA-CFA expires.
B. If the virtual interface is not created within 72 days, the LOA-CFA becomes outdated.

C. If the cross connect is not completed within a user-defined time, the authority granted by the LOA- CFA expires.
D. If the cross connect is not completed within the specified duration from the appropriate provider, the LOA-CFA expires.
Answer: (SHOW ANSWER)
An AWS Direct Connect location provides access to AWS in the region it is associated with. You can establish connections with AWS Direct Connect locations in multiple regions, but a connection in one region does not provide connectivity to other regions. Note: If the cross connect is not completed within 90 days, the authority granted by the LOA-CFA expires.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Colocation.html

NEW QUESTION: 35
A user has created a VPC with public and private subnets using the VPC Wizard. The VPC has CIDR 20.0.0.0/16. The private subnet uses CIDR 20.0.0.0/24.
Which of the below mentioned entries are required in the main route table to allow the instances in VPC to communicate with each other?
A. Destination : 20.0.0.0/0 and Target : ALL
B. Destination : 20.0.0.0/16 and Target : Local
C. Destination : 20.0.0.0/24 and Target : Local
D. Destination : 20.0.0.0/16 and Target : ALL
Answer: B (LEAVE A REPLY)
Explanation
A user can create a subnet with VPC and launch instances inside that subnet. If the user has created a public private subnet, the instances in the public subnet can receive inbound traffic directly from the Internet, whereas the instances in the private subnet cannot. If these subnets are created with Wizard, AWS will create two route tables and attach to the subnets. The main route table will have the entry "Destination: 20.0.0.0/16 and Target: Local", which allows all instances in the VPC to communicate with each other.
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html

NEW QUESTION: 36
The CFO of a company wants to allow one of his employees to view only the AWS usage report page.
Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?
A. "Effect": "Allow", "Action": ["Describe"], "Resource": "Billing"
B. "Effect": "Allow", "Action": ["aws-portal: ViewBilling"], "Resource": "*"
C. "Effect": "Allow", "Action": ["aws-portal: ViewUsage"], "Resource": "*"
D. "Effect": "Allow", "Action": ["AccountUsage], "Resource": "*"
Answer: C (LEAVE A REPLY)
Explanation

AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow", "Action": [
"aws-portal:ViewUsage"
],
"Resource": "*"
}
]
}
```

http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html

NEW QUESTION: 37

A Solutions Architect must design a storage solution for incoming billing reports in CSV format. The data does not need to be scanned frequently and is discarded after 30 days.

Which service will be MOST cost-effective in meeting these requirements?

A. Import the logs to an Amazon Redshift cluster

B. Write the files to an S3 bucket and use Amazon Athena to query the data.

C. Use AWS Data Pipeline to import the logs into a DynamoDB table.

D. Import the logs into an RDS MySQL instance.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 38

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose ones.

streqi is the short version of the _____ string condition.

A. StringEqualsIgnoreCase

B. StringNotEqualsIgnoreCase

C. StringLikeStringEquals

D. StringNotEquals

Answer: A (LEAVE A REPLY)

When using string conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, streqi is the short version of StringEqualsIgnoreCase that checks for the exact match between two strings ignoring their case.

http://awsdocs.s3.amazonaws.com/SNS/20100331/sns-gsg-2010-03-31.pdf

NEW QUESTION: 39

Pass4test has created a multi-tenant Learning Management System (LMS). The application is hosted for five different tenants (clients) in the VPCs of the respective AWS accounts of the tenant. Pass4test wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required.
Pass4test also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons. How can Pass4test setup this scenario?
A. Pass4test has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
B. Pass4test should setup VPC peering with all the VPCs peering each other but block the IPs from CIDR of the tenant VPCs to deny them.
C. Pass4test should setup all the VPCs with the same CIDR but have a centralized VPC. This way only the centralized VPC can talk to the other VPCs using VPC peering.
D. Pass4test should setup all the VPCs meshed together with VPC peering for all VPCs.
Answer: A (LEAVE A REPLY)
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network. This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC. The organization wants to setup that one VPC can connect with all the other VPCs but all other VPCs cannot connect among each other. This can be achieved by configuring VPC peering where one VPC is peered with all the other VPCs, but the other VPCs are not peered to each other. The VPCs are in the same or a separate AWS account and should not have overlapping CIDR blocks.
http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html#many-vpcs-full-acces

NEW QUESTION: 40
A prediction process requires access to a trained model that is stored in an Amazon S3 bucket. The process takes a few seconds to process an image and make a prediction. The process takes a few seconds to process an image and make a prediction. The process is not overly resource-intensive, does not require any specialized hardware, and takes less than 512 MB of memory to run.
What would be the MOST effective compute solution for this use case?
A. Amazon EC2 Spot instances
B. AWS Lambda functions
C. AWS Elastic Beanstalk
D. Amazon ECS
Answer: B (LEAVE A REPLY)

NEW QUESTION: 41
A company runs a video processing platform. Files are uploaded by users who connect to a web server, which stores them on an Amazon EFS share. This web server is running on a single Amazon EC2 instance. A different group of instances, running in an Auto Scaling group, scans the EFS share directory structure for new files to process and generates new videos (thumbnails, different resolution, compression, etc.) according to the instructions file, which is uploaded along with the video files. A different application running on a group of

instances managed by an Auto Scaling group processes the video files and then deletes them from the EFS share. The results are stored in an S3 bucket. Links to the processed video files are emailed to the customer. The company has recently discovered that as they add more instances to the Auto Scaling Group, many files are processed twice, so image processing speed is not improved. The maximum size of these video files is 2GB.
What should the Solutions Architect do to improve reliability and reduce the redundant processing of video files?
A. Modify the web application to upload the video files directly to Amazon S3. Use Amazon CloudWatch Events to trigger an AWS Lambda function every time a file is uploaded, and have this Lambda function put a message into an Amazon queue for new files and use the queue depth metric to scale instances in the video processing Auto Scaling group.
B. Set up a cron job on the web server instance to synchronize the contents of the EFS share into Amazon S3. Trigger an AWS Lambda function every time a file is uploaded to process the video file and store the results in Amazon S3. Using Amazon CloudWatch Events trigger an Amazon SES job to send an email to the customer containing the link to the processed file.
C. Rewrite the web application to run directly from Amazon S3 and use Amazon API Gateway to upload the video files to an S3 bucket. Use an S3 trigger to run an AWS Lambda function each time a file is uploaded to process and store new video files in a different bucket. Using CloudWatch Events, trigger an SES job to send an email to the customer containing the link to the processed file.
D. Rewrite the application to run from Amazon S3 and upload the video files to an S3 bucket. Each time a new file is uploaded, trigger an AWS Lambda function to put a message in an SQS queue containing the link and the instructions. Modify the video processing application to read from the SQS queue and the S3 bucket. Use the queue depth metric to adjust the size of the Auto Scaling group for video processing instances.
Answer: C (LEAVE A REPLY)
A\D: If SQS is used, then FIFO must be configured to ensure that the messages are processed only once. https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO- queues.html B: There is a time lag in the sync from EFS to S3.
C: This is more instant compared to A and hence improved the reliability.

NEW QUESTION: 42
A finance company is running its business-critical application on current-generation Linux EC2 instances.
The application includes a self-managed MySQL database performing heavy I/O operations. The application is working fine to handle a moderate amount of traffic during the month. However, it slows down during the final three days of each month due to month-end reporting, even though the company is using Elastic Load Balancers and Auto Scaling within its infrastructure to meet the increased demand.
Which of the following actions would allow the database to handle the month-end load with the LEAST impact on performance?
A. Performing a one-time migration of the database cluster to Amazon RDS, and creating several additional read replicas to handle the load during end of month.
B. Using Amazon CloudWatch with AWS Lambda to change the type, size, or IOPS of Amazon EBS volumes in the cluster based on a specific CloudWatch metric.
C. Replacing all existing Amazon EBS volumes with new PIOPS volumes that have the maximum available storage size and I/O per second by taking snapshots before the end of the month and reverting back afterwards.

D. Pre-warming Elastic Load Balancers, using a bigger instance type, changing all Amazon EBS volumes to GP2 volumes.
Answer: (SHOW ANSWER)

NEW QUESTION: 43
A company is planning to host a web application on AWS and works to load balance the traffic across a group of Amazon EC2 instances. One of the security requirements is to enable end-to-end encryption in transit between the client and the web server.
Which solution will meet this requirement?
A. Place the EC2 instances behind an Application Load Balancer (ALB) Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Export the SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
B. Associate the EC2 instances with a target group. Provision an SSL certificate using AWS Certificate Manager (ACM). Create an Amazon CloudFront distribution and configure It to use the SSL certificate.
Set CloudFront to use the target group as the origin server
C. Place the EC2 instances behind an Application Load Balancer (ALB). Provision an SSL certificate using AWS Certificate Manager (ACM), and associate the SSL certificate with the ALB. Provision a third-party SSL certificate and install it on each EC2 instance. Configure the ALB to listen on port 443 and to forward traffic to port 443 on the instances.
D. Place the EC2 instances behind a Network Load Balancer (NLB). Provision a third-party SSL certificate and install it on the NLB and on each EC2 instance. Configure the NLB to listen on port 443 and to forward traffic to port 443 on the instances.
Answer: (SHOW ANSWER)
Explanation
Using an Application Load Balancer (ALB) to distribute traffic to the EC2 instances is the best solution to meet the requirement of enabling end-to-end encryption in transit between the client and the web server. The ALB should be configured to listen on port 443 and forward traffic to port 443 on the EC2 instances. An SSL certificate should be provisioned using AWS Certificate Manager (ACM) and associated with the ALB. The SSL certificate should then be exported and installed on each EC2 instance.
* AWS Application Load Balancer documentation:
https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/
* AWS Certificate Manager documentation: https://aws.amazon.com/certificate-manager/
* AWS Encryption in Transit documentation:
https://aws.amazon.com/premiumsupport/knowledge-center/encryption-in-transit-alb-nlb/

NEW QUESTION: 44
A company is planning on hosting its ecommerce platform on AWS using a multi-tier web application designed for a NoSQL database. The company plans to use the us-west-2 Region as its primary Region. The company want to ensure that copies of the application and data are available in second Region, us-west-1, for disaster

recovery. The company wants to keep the time to fall over as tow as possible. Fading back to the primary Region should be possible without administrative interaction after the primary service is restored.
Which design should the solutions architect use?
A. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tiers Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage Use Amazon DynamoOB global tables for the database tier
B. Use AWS CloudFormation StackSets to create the stacks in both Regions with Auto Scaling groups for the web and application tiers Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication. Use an Amazon Route 53 DNS failover routing policy to direct users to the secondary site in us-west-1 in the event of an outage Deploy an Amazon Aurora global database for the database tier
C. Use AWS Service Catalog to deploy the web and application servers in both Regions Asynchronously replicate static content between the two Regions using Amazon S3 cross-Region replication Use Amazon Route 53 health checks to identify a primary Region failure and update the public DNS entry listing to the secondary Region in the event of an outage. Use Amazon RDS for MySQL with cross-Region replication for the database tier
D. Use AWS CloudFormation StackSets to create the stacks in both Regions using Auto Scaling groups for the web and application bars Asynchronously replicate static content between Regions using Amazon S3 cross-Region replication Use Amazon CloudFront with static files in Amazon S3, and multi-Region origins tor the front-end web tier. Use Amazon DynamoDB tables m each Region with scheduled backups to Amazon S3
Answer: A (LEAVE A REPLY)

NEW QUESTION: 45
A company is using AWS to run an internet-facing production application written in Node.js. The Development team is responsible for pushing new versions of their software directly to production. The application software is updated multiple times a day. The team needs guidance from a Solutions Architect to help them deploy the software to the production fleet quickly and with the least amount of disruption to the service.
Which option meets these requirements?
A. Deploy the base AMI through Auto Scaling and bootstrap the software using user data. For software changes, SSH to each of the instances and replace the software with the new version.
B. Use AWS CodeDeploy to push the prepackaged AMI to production. For software changes, reconfigure CodeDeploy with new AMI identification to push the new AMI to the production fleet.
C. Prepackage the software into an AMI and then use Auto Scaling to deploy the production fleet. For software changes, update the AMI and allow Auto Scaling to automatically push the new AMI to production.
D. Use AWS Elastic Beanstalk to host the production application. For software changes, upload the new application version to Elastic Beanstalk to push this to the production fleet using a blue/green deployment method.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 46

A company wants to retire its Oracle Solaris NFS storage arrays. The company requires rapid data migration over its internet network connection to a combination of destinations for Amazon S3. Amazon Elastic File System (Amazon EFS), and Amazon FSx lor Windows File Server. The company also requires a full initial copy, as well as incremental transfers of changes until the retirement of the storage arrays. All data must be encrypted and checked for integrity.
What should a solutions architect recommend to meet these requirements?
A. Configure CloudEndure. Create a project and deploy the CloudEndure agent and token to the storage array. Run the migration plan to start the transfer.
B. Configure AWS DataSync. Configure the DataSync agent and deploy it to the local network. Create a transfer task and start the transfer.
C. Configure the aws S3 sync command. Configure the AWS client on the client side with credentials. Run the sync command to start the transfer.
D. Configure AWS Transfer for FTP. Configure the FTP client with credentials. Script the client to connect and sync to start the transfer.
Answer: B (LEAVE A REPLY)
Explanation
It enables secure, high-performance transfers and supports both full initial copies and incremental transfers of changes. DataSync provides encryption and checksum validation to ensure data integrity, and it can be configured to transfer data over the internet or over a private network connection. Additionally, it can be scripted and automated, making it a great choice for this scenario.
References:
AWS Documentation - DataSync
https://docs.aws.amazon.com/datasync/latest/userguide/what-is-datasync.html AWS Certified Solutions Architect Professional Official Text Book - Chapter 4. Data Storage
https://aws.amazon.com/training/learning-paths/professional-solutions-architect/

NEW QUESTION: 47
A company's web application is running on Amazon EC2 instances behind an Application Load Balancer.
The company recently changed its policy, which now requires the application to be accessed from one specific country only.
Which configuration will meet this requirement?

A. Configure the security group for the EC2 instances.
B. Configure the security group on the Application Load Balancer.
C. Configure AWS WAF on the Application Load Balancer in a VPC.
D. Configure the network ACL for the subnet that contains the EC2 instances.
Answer: C (LEAVE A REPLY)
Explanation/Reference:
https://aws.amazon.com/es/blogs/security/how-to-use-aws-waf-to-filter-incoming-traffic-from-embargoed-countries/

NEW QUESTION: 48
A startup company hosts a fleet of Amazon EC2 instances in private subnets using the latest Amazon Linux 2 AMI The company's engineers rely heavily on SSH access to the instances for troubleshooting.
The company's existing architecture includes the following
* A VPC with private and public subnets, and a NAT gateway
* Site-to-Site VPN for connectivity with the on-premises environment
* EC2 security groups with direct SSH access from the on-premises environment The company needs to increase security controls around SSH access and provide auditing of commands executed by the engineers
Which strategy should a solutions architect use"
A. Install and configure EC2 Instance Connect on the fleet of EC2 instances Remove all security group rules attached to EC2 instances that allow Inbound TCP on port 22 Advise the engineers to remotely access the instances by using the EC2 Instance Connect CLI
B. Update the EC2 security groups to only allow Inbound TCP on port 22 to the IP addresses of the engineer's devices Install the Amazon CloudWatch agent on all EC2 instances and send operating system audit logs to CloudWatch Logs
C. Create an 1AM role with the AmazonSSMManaged InstanceCore managed policy attached Attach the 1AM role to all the EC2 instances Remove all security group rules attached to the EC2 instances that allow inbound TCP on port 22. Have the engineers install the AWS Systems Manager Session Manager plugin lor their devices and remotely access the instances by using the start-session API call from Systems Manager
D. Update the EC2 security groups to only allow Inbound TCP on port 22 to the IP addresses of the engineer's devices Enable AWS Config for EC2 security group resource changes Enable AWS Firewall Manager and apply a security group policy that automatically remediates changes to rules
Answer: (SHOW ANSWER)

NEW QUESTION: 49
Which of the following statements is NOT correct when working with your AWS Direct Connect connection after it is set up completely?
A. You can manage your AWS Direct Connect connections and view the connection details.
B. You can delete a connection as long as there are no virtual interfaces attached to it.
C. You cannot view the current connection ID and verify if it matches the connection ID on the Letter of Authorization (LOA).
D. You can accept a host connection by purchasing a hosted connection from the partner (APN).

Answer: C (LEAVE A REPLY)

Explanation

You can manage your AWS Direct Connect connections and view connection details, accept hosted connections, and delete connections. You can view the current status of your connection. You can also view your connection ID, which looks similar to this example dxcon-xxxx, and verify that it matches the connection ID on the Letter of Authorization (LOA) that you received from Amazon.

http://docs.aws.amazon.com/directconnect/latest/UserGuide/viewdetails.html

NEW QUESTION: 50

A company has a project that is launching Amazon EC2 instances that are larger than required. The project's account cannot be part of the company's organization in AWS Organizations due to policy restrictions to keep this activity outside of corporate IT. The company wants to allow only the launch of t3.small EC2 instances by developers in the project's account. These EC2 instances must be restricted to the us-east-2 Region.

What should a solutions architect do to meet these requirements?

A. Create a new developer account. Move all EC2 instances, users, and assets into us-east-2. Add the account to the company's organization in AWS Organizations. Enforce a tagging policy that denotes Region affinity.

B. Create an SCP that denies the launch of all EC2 instances except I3.small EC2 instances in us-east-2. Attach the SCP to the project's account.

C. Create and purchase a t3.small EC2 Reserved Instance for each developer in us-east-2. Assign each developer a specific EC2 instance with their name as the tag.

D. Create an 1AM policy than allows the launch of only t3.small EC2 instances in us-east-2. Attach the policy to the roles and groups that the developers use in the project's account.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 51

A company used Amazon EC2 instances to deploy a web fleet to host a blog site. The EC2 instances are behind an Application Load Balancer (ALB) and are configured in an Auto Scaling group. The web application stores all blog content on an Amazon EFS volume.

The company recently added a feature for bloggers to add video to their posts, attracting 10 times the previous user traffic. At peak times of day, users report buffering and timeout issues while attempting to reach the site or watch videos.

Which is the MOST cost-efficient and scalable deployment that will resolve the issues for users?

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

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

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

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

Answer: C (LEAVE A REPLY)

A: Issue seems to be latency and load related. EFS does not solve the issue since the issue lies with EC2.

B: Risky as an EC2 instance failure could corrupt the data.

D: Origin cannot point to ALB (either S3, EC2 or HTTP based)?

NEW QUESTION: 52
A health insurance company stores personally identifiable information (PII) in an Amazon S3 bucket. The company uses server-side encryption with S3 managed encryption keys (SSE-S3) to encrypt the objects. According to a new requirement, all current and future objects in the S3 bucket must be encrypted by keys that the company's security team manages. The S3 bucket does not have versioning enabled.
Which solution will meet these requirements?
A. In the S3 bucket properties, change the default encryption to SSE-S3 with a customer managed key. Use the AWS CLI to re-upload all objects in the S3 bucket. Set an S3 bucket policy to deny unencrypted PutObject requests.
B. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to automatically encrypt objects on GetObject and PutObject requests.
C. In the S3 bucket properties, change the default encryption to AES-256 with a customer managed key. Attach a policy to deny unencrypted PutObject requests to any entities that access the S3 bucket. Use the AWS CLI to re-upload all objects in the S3 bucket.
D. In the S3 bucket properties, change the default encryption to server-side encryption with AWS KMS managed encryption keys (SSE-KMS). Set an S3 bucket policy to deny unencrypted PutObject requests.
Use the AWS CLI to re-upload all objects in the S3 bucket.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 53
A start-up company has a web application based in the us-east-1 Region with multiple Amazon EC2 instances running behind an Application Load Balancer across multiple Availability Zones. As the company's user base grows in the us-west-1 Region, it needs a solution with low latency and high availability.
What should a solutions architect do to accomplish this?
A. Provision EC2 instances and configure an Application Load Balancer in us-west-1. Configure Amazon Route 53 with a weighted routing policy. Create alias records in Route 53 that point to the Application Load Balancer.
B. Provision EC2 instances and configure an Application Load Balancer in us-west-1. Create an accelerator in AWS Global Accelerator that uses an endpoint group that includes the load balancer endpoints in both Regions.
C. Provision EC2 instances in us-west-1. Switch the Application Load Balancer to a Network Load Balancer to achieve cross-Region load balancing.
D. Provision EC2 instances and an Application Load Balancer in us-west-1. Make the load balancer distribute the traffic based on the location of the request.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 54
Your team has a tomcat-based java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a

sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The optimal setup for persistence and security that meets the above requirements would be the following:
A. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code.
Alter its security group to allow access to it from hosts within your VPC's IP address block.
B. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable.
Alter its security group to allow access to it from hosts in your application subnets.
C. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
D. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable.
Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
Answer: D (LEAVE A REPLY)
Elastic Beanstalk provides support for running Amazon RDS instances in your Elastic Beanstalk environment. This works great for development and testing environments, but is not ideal for a production environment because it ties the lifecycle of the database instance to the lifecycle of your application's environment.
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html

NEW QUESTION: 55
The Security team needs to provide a team of interns with an AWS environment so they can build the serverless video transcoding application. The project will use Amazon S3, AWS Lambda, Amazon API Gateway, Amazon Cognito, Amazon DynamoDB, and Amazon Elastic Transcoder.
The interns should be able to create and configure the necessary resources, but they may not have access to create or modify AWS IAM roles. The Solutions Architect creates a policy and attaches it to the interns' group.
How should the Security team configure the environment to ensure that the interns are self-sufficient?
A. Create a policy that allows creation of all project-related resources, including roles that allow access only to specified resources.
B. Create a policy that allows creation of project-related resources only. Require the interns to raise a request for roles to be created with the Security team. The interns will provide the requirements for the permissions to be set in the role.
C. Create a policy that allows creation of project-related resources only. Create roles with required service permissions, which are assumable by the services.
D. Create roles with the required service permissions, which are assumable by the services. Have the interns create and use a bastion host to create the project resources in the project subnet only.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 56

A North American company with headquarters on the East Coast is deploying a new web application running on Amazon EC2 in the us-east-1 Region. The application should dynamically scale to meet user demand and maintain resiliency. Additionally, the application must have disaster recover capabilities in an active-passive configuration with the us-west-1 Region.
Which steps should a solutions architect take after creating a VPC in the us-east-1 Region?
A. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region. Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALB Deploy the same solution to the us-west-1 Region Create separate Amazon Route 53 records in each Region that point to the ALB in the Region. Use Route 53 health checks to provide high availability across both Regions.
B. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region Deploy EC2 instances across multiple AZs as part of an Auto Scaling group served by the ALB. Deploy the same solution to the us-west-1 Region. Create an Amazon Route 53 record set with a failover routing policy and health checks enabled to provide high availability across both Regions.
C. Create a VPC in the us-west-1 Region Use inter-Region VPC peering to connect both VPCs. Deploy an Application Load Balancer (ALB) spanning multiple Availability Zones (AZs) to the VPC in the us-east-1 Region Deploy EC2 instances across multiple AZs in each Region as part of an Auto Scaling group spanning both VPCs and served by the ALB.
D. Create a VPC in the us-west-1 Region Use inter-Region VPC peering to connect both VPCs. Deploy an Application Load Balancer (ALB) that spans both VPCs. Deploy LC2 instances across multiple Availability Zones as part of an Auto Scaling group in each VPC served by the ALB. Create an Amazon Route 53 record that points to the ALB.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 57
What bandwidths do AWS Direct Connect currently support?
A. 10Mbps and 100Mbps
B. 10Gbps and 100Gbps
C. 100Mbps and 1Gbps
D. 1Gbps and 10 Gbps
Answer: D (LEAVE A REPLY)
Explanation
AWS Direct Connection currently supports 1Gbps and 10 Gbps.
http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

NEW QUESTION: 58
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: (SHOW ANSWER)
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: 59
A company is migrating a legacy application from an on-premises data center to AWS. The application uses MangeDB as a key-value database According to the company's technical guidelines, all Amazon EC2 instances must be hosted in a private subnet without an internet connection In addition, all connectivity between applications and databases must be encrypted. The database must be able to scale based on demand Which solution will meet these requirements?
A. Create new Amazon DocumentDB (with MangeDB compatibility) tables for the application with Provisioned IOPS volumes Use the instance endpoint to connect to Amazon DocumentDB

B. Create new Amazon DynamoDB tables for the application with on-demand capacity Use a gateway VPC endpoint for DynamoDB to connect lo the DynamoDB tables
C. Create new Amazon DynamoDB tables for the application with on-demand capacity Use an interface VPC endpoint for DynamoDB to connect to the DynamoDB tables
D. Create new Amazon DocumentDB (with MangeDB compatibility) tables for the application with Provisioned IOPS volumes Use the cluster endpoint to connect to Amazon DocumentDB
Answer: B (LEAVE A REPLY)
Explanation
https://docs.aws.amazon.com/vpc/latest/privatelink/aws-services-privatelink-support.html

NEW QUESTION: 60
A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public facing ELB. Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.
How should they architect their solution?
A. Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the MAT instances.
B. Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instances public IP address to the payment validation whitelist API.
C. Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.
D. Whitelist the ELB IP addresses and route payment requests from the Application servers through the ELB.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 61
An AWS partner company is building a service in AWS Organizations using its organization named org1.
This service requires the partner company to have access to AWS resources in a customer account, which is in a separate organization named org2. The company must establish least privilege security access using an API or command line tool to the customer account.
What is the MOST secure way to allow org1 to access resources in org2?
A. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) when requesting access to perform the required tasks.
B. The customer should provide the partner company with their AWS account access keys to log in and perform the required tasks.
C. The customer should create an IAM role and assign the required permissions to the IAM role. The partner company should then use the IAM role's Amazon Resource Name (ARN) including the external ID in the IAM role's trust policy when requesting access to perform the required tasks.

D. 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.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 62
To ensure failover capabilities on an elastic network interface (ENI), what should you use for incoming traffic?
A. A Route53 A record
B. A secondary private IP
C. A secondary public IP
D. A secondary ENI
Answer: B (LEAVE A REPLY)
Explanation
To ensure failover capabilities on an elastic network interface (ENI), consider using a secondary private IP for incoming traffic and if a failure occurs, you can move the interface and/or secondary private IP address to a standby instance.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

NEW QUESTION: 63
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 privilege security access using an API or command line tool to the customer account What is the MOST secure way to allow org1 to access resources 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)

Explanation

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: 64

A company Is serving files to its customers through an SFTP server that Is accessible over the internet The SFTP server Is running on a single Amazon EC2 instance with an Elastic IP address attached Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication The EC2 instance also has an attached security group that allows access from all customer IP addresses.

A solutions architect must implement a solution to improve availability minimize the complexity ot infrastructure management and minimize the disruption to customers who access files. The solution must not change the way customers connect.

Which solution will meet these requirements?

A. Disassociate the Elastic IP address from me EC2 instance Create an Amazon S3 bucket to be used for sftp file hosting Create an AWS Transfer Family server Configure the Transfer Family server with a publicly accessible endpoint. Associate the SFTP Elastic IP address with the new endpoint. Point the Transfer Family server to the S3 bucket Sync all files from the SFTP server to the S3 bucket.

B. Disassociate the Elastic IP address from the EC2 instance. Create an Amazon S3 bucket to be used for SFTP file hosting Create an AWS Transfer Family server. Configure the Transfer Family server with a VPC-hosted. internet-facing endpoint. Associate the SFTP Elastic IP address with the new endpoint.

Attach the security group with customer IP addresses to the new endpoint. Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to The S3 bucket

C. Disassociate the Elastic IP address from the EC2 instance. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting. Create an AWS Fargate task definition to run an SFTP server. Specify the EFS file system as a mount in the task definition Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) front of the service When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP server Associate the Elastic IP address with the Nl B Sync all files from the SFTP server to the S3 bucket

D. Disassociate the Elastic IP address from the EC2 instance Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used to SFTP file hosting Create a Network Load Balancer (NLB) with the Elastic IP address attached Create an Auto Scaling group with EC2 instances that run an SFTP server Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume Configure the Auto Scaling group to automatically add instances behind the NLB Configure the Auto Scaling group to use the

security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches
Sync all files from the SFTP server to the new multi-attach EBS volume
Answer: (SHOW ANSWER)
Explanation
https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/
https://docs.aws.amazon.com/transfer/latest/userguide/create-server-in-vpc.html
https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/

NEW QUESTION: 65
A Solutions Architect has created an AWS CloudFormation template for a three-tier application that contains an Auto Scaling group of Amazon EC2 instances running a custom AMI.
The Solutions Architect wants to ensure that future updates to the custom AMI can be deployed to a running stack by first updating the template to refer to the new AMI, and then invoking UpdateStack to replace the EC2 instances with instances launched from the new AMI.
How can updates to the AMI be deployed to meet these requirements?
A. Create a change set for a new version of the template, view the changes to the running EC2 instances to ensure that the AMI is correctly updated, and then execute the change set.
B. Edit the AWS::AutoScaling::LaunchConfiguration resource in the template, changing its DeletionPolicy to Replace.
C. Edit the AWS::AutoScaling:: AutoScalingGroup resource in the template, inserting an UpdatePolicy attribute.
D. Create a new stack from the updated template. Once it is successfully deployed, modify the DNS records to point to the new stack and delete the old stack.
Answer: C (LEAVE A REPLY)
Explanation
References:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-as-launchconfig.html

NEW QUESTION: 66
A company has a web service deployed in the following two AWS Regions: us-west-2 and us-est-1. Each AWS region runs an identical version of the web service. Amazon Route 53 is used to route customers to the AWS Region that has the lowest latency.
The company wants to improve the availability of the web service in case an outage occurs in one of the two AWS Regions.
A Solutions Architect has recommended that a Route 53 health check be performed. The health check must detect a specific text on an endpoint.
What combination of conditions should the endpoint meet to pass the Route 53 health check? (Choose two.)
A. The endpoint must establish a TCP connection within 10 seconds.
B. The specific text string must appear within the first 5,120 bytes of the response.
C. The endpoint must return an HTTP 200 status code.
D. The endpoint must respond to the request within the number of seconds specified when creating the health check.

E. The endpoint must return an HTTP 2xx or 3xx status code.
Answer: (SHOW ANSWER)

NEW QUESTION: 67
A company has a data center that must be migrated to AWS as quickly as possible. The data center has a 500 Mbps AWS Direct Connect link and a separate, fully available 1 Gbps ISP connection. A Solutions Architect must transfer 20 TB of data from the data center to an Amazon S3 bucket.
What is the FASTEST way transfer the data?
A. Upload the data to the S3 bucket using the existing DX link.
B. Send the data to AWS using the AWS Import/Export service.
C. Upload the data using an 80 TB AWS Snowball device.
D. Upload the data to the S3 bucket using S3 Transfer Acceleration.
Answer: D (LEAVE A REPLY)
S3 Transfer Acceleration enables fast, easy, and secure transfers of files over long distances between your client and an S3 bucket, and it takes about 2 days to upload 20TB data via 1Gbps ISP connection.

NEW QUESTION: 68
An ERP application is deployed across multiple AZs in a single region. In the event of failure, the Recovery Time Objective (RTO) must be less than 3 hours, and the Recovery Point Objective (RPO) must be 15 minutes. The customer realizes that data corruption occurred roughly 1.5 hours ago.
What DR strategy could be used to achieve this RTO and RPO in the event of this kind of failure?
A. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored In S3 every 5 minutes.
B. Take hourly DB backups to S3, with transaction logs stored in S3 every 5 minutes.
C. Take 15 minute DB backups stored In Glacier with transaction logs stored in S3 every 5 minutes.
D. Use synchronous database master-slave replication between two availability zones.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 69
The Solutions Architect manages a serverless application that consists of multiple API gateways, AWS Lambda functions, Amazon S3 buckets, and Amazon DynamoDB tables. Customers say that a few application components slow while loading dynamic images, and some are timing out with the "504 Gateway Timeout" error. While troubleshooting the scenario, the Solutions Architect confirms that DynamoDB monitoring metrics are at acceptable levels.
Which of the following steps would be optimal for debugging these application issues? (Choose two.)
A. Parse HTTP logs in Amazon API Gateway for HTTP errors to determine the root cause of the errors.
B. Parse Amazon CloudWatch Logs to determine processing times for requested images at specified intervals.
C. Parse VPC Flow Logs to determine if there is packet loss between the Lambda function and S3.
D. Parse AWS X-Ray traces and analyze HTTP methods to determine the root cause of the HTTP errors.
E. Parse S3 access logs to determine if objects being accessed are from specific IP addresses to narrow the scope to geographic latency issues.
Answer: A,D (LEAVE A REPLY)

A: https://aws.amazon.com/about-aws/whats-new/2017/11/amazon-api-gateway-supports- access-logging/ B: https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/http-504-gateway- timeout.html >> 504 already indicates problem at the origin, why do we want to still determine the processing time?
C: This is not useful as 504 error points to timeout between the frontend and the backend (origin).
D: https://docs.aws.amazon.com/en_pv/apigateway/latest/developerguide/apigateway-xray.html E: This is not a latency issue. It's the Cloudfront unable to download from origin.

NEW QUESTION: 70
You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet.
The application must have a highly available architecture.
Which alternatives should you consider? (Choose 2 answers)
A. Configure ELB with an EIP Place all your Web servers behind ELB Configure a Route53 A record that points to the EIP.
B. Place all your web servers behind ELB Configure a Route53 CNMIE to point to the ELB DNS name.
C. Configure a CloudFront distribution and configure the origin to point to the private IP addresses of your Web servers Configure a Route53 CNAME record to your CloudFront distribution.
D. Configure a NAT instance in your VPC Create a default route via the NAT instance and associate itwith all subnets Configure a DNS A record that points to the NAT instance public IP address.
E. Assign EIPs to all web servers. Configure a Route53 record set with all EIPs, with health checks and DNS failover.
Answer: B,E (LEAVE A REPLY)

NEW QUESTION: 71
A company is planning to host a three tier application in the AWS Cloud The application layer will use Amazon EC2 in an Auto Scaling group A custom EC2 role named AppServer will be created and associated with the application instances The entire application stack will be deployed using AWS Cloud Formation The company's security team requires encryption of all AMI snapshots and Amazon Plastic Block Store (Amazon TBS) volumes with an AWS Key Management Service (AWS KMS> CMK Which action will deploy the stack correctly after the AMI snapshot is encrypted with the KMS key?
A. Update the KMS key policy to provide the required permissions to the AWSServiceRoleForAutoScalir>g service-linked role
B. Update the AppServer role to have the required permissions to access the KMS key
C. Update the KMS key policy to provide the required permissions to the AppServer role
D. Update the CloudFormation stack role to have the required permissions to access the KMS key
Answer: (SHOW ANSWER)

NEW QUESTION: 72
In DynamoDB, a projection is _____.
A. systematic transformation of the latitudes and longitudes of the locations inside your table
B. importing data from your file to a table
C. exporting data from a table to your file

D. the set of attributes that is copied from a table into a secondary index
Answer: (SHOW ANSWER)
In DynamoDB, a projection is the set of attributes that is copied from a table into a secondary index.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GSI.html

NEW QUESTION: 73
An AWS customer is deploying an application mat is composed of an AutoScaling group of EC2 Instances.
The customers security policy requires that every outbound connection from these instances to any other service within the customers Virtual Private Cloud must be authenticated using a unique x 509 certificate that contains the specific instance-id.
In addition, an x 509 certificates must Designed by the customer's Key management service in order to be trusted for authentication.
Which of the following configurations will support these requirements?
A. Embed a certificate into the Amazon Machine Image that is used by the Auto Scaling group. Have the launched instances generate a certificate signature request with the instance's assigned instance-id to the key management service for signature.
B. Configure an IAM Role that grants access to an Amazon S3 object containing a signed certificate and configure the Auto Scaling group to launch instances with this role. Have the instances bootstrap get the certificate from Amazon S3 upon first boot.
C. Configure the launched instances to generate a new certificate upon first boot. Have the Key management service poll the Auto Scaling group for associated instances and send new instances a certificate signature (hat contains the specific instance-id.
D. Configure the Auto Scaling group to send an SNS notification of the launch of a new instance to the trusted key management service. Have the Key management service generate a signed certificate and send it directly to the newly launched instance.
Answer: (SHOW ANSWER)

NEW QUESTION: 74
A solutions architect is auditing the security setup of an AWS Lambda function for a company. The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.
The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.
What should the solutions architect recommend to meet these requirements?
A. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

B. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Enforce HTTPS on the connection to Amazon S3 during data transfers.
C. Save the database credentials in AWS Systems Manager Parameter Store. Set up password rotation on the credentials in Parameter Store. Change the IAM role for the Lambda function to allow the function to access Parameter Store. Modify the Lambda function to retrieve the credentials from Parameter Store.
Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
D. Save the database credentials in AWS Secrets Manager. Set up password rotation on the credentials in Secrets Manager. Change the IAM role for the Lambda function to allow the function to access Secrets Manager. Modify the Lambda function to retrieve the credentials Om Secrets Manager. Enforce HTTPS on the connection to Amazon S3 during data transfers.
Answer: A (LEAVE A REPLY)
Explanation
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.html

NEW QUESTION: 75
A company wants to improve cost awareness for its Amazon EMR platform. The company has allocated budgets for each team's Amazon EMR usage. When a budgetary threshold is reached, a notification should be sent by email to the budget office's distribution list. Teams should be able to view their EMR cluster expenses to date. A solutions architect needs to create a solution that ensures the policy is proactrvely and centrally enforced in a multi-account environment.
Which combination of steps should the solutions architect take to meet these requirements? (Select TWO)
A. Implement Amazon CloudWatch dashboards for Amazon EMR usage
B. Create an AWS Service Catalog portfolio tor each team. Add each team's Amazon EMR cluster as an AWS Cloud Formationtemplate to their Service Catalog portfolio as a Product.
C. Update the AWS ClouddFormation template to include the AWS::Budgets::Budget::resource with the NotificationsWithSubscribers property.
D. Create an EMR bootstrap action that runs at startup that calls the Cost Explorer API to set the budget on the cluster with the GetCostForecast and NotificationsWithSubscribers actions.
E. Create an Amazon CloudWatch metric for billing. Create a custom alert when costs exceed the budgetary threshold.
Answer: B,E (LEAVE A REPLY)

NEW QUESTION: 76
A company needs to create and manage multiple AWS accounts for a number ol 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 tor the security team.
How should a solutions architect meet these requirements?
A. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole 1AM role in the member account from the security account. Use the generated temporary credentials to gain access.

B. Use the OrganizationAccountAccessRole 1AM role to create a new 1AM policy wilh read-only access in each member account. Establish a trust relationship between the 1AM policy in each member account and the security account. Ask the security team lo use the 1AM policy to gain access.
C. Use the OrganizationAccountAccessRole 1AM role to create a new 1AM role with read-only access in each member account. Establish a trust relationship between the 1AM role in each member account and the security account. Ask the security team lo use the 1AM role to gain access.
D. Ask the security team to use AWS Security Token Service (AWS STS) to call the AssumeRole API for the OrganizationAccountAccessRole 1AM role in the master account from the security account. Use the generated temporary credentials to gain access.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 77
A company is hosting multiple websites for several lines of business under its registered parent domain. Users accessing these websites will be routed to appropriate backend Amazon EC2 instances based on the subdomain. The websites host static webpages, images, and server-side scripts like PHP and JavaScript. Some of the websites experience peak access during the first two hours of business with constant usage throughout the rest of the day. A solutions architect needs to design a solution that will automatically adjust capacity to these traffic patterns while keeping costs low.
Which combination of AWS services or features will meet these requirements? (Choose two.)
A. Application Load Balancer
B. Network Load Balancer
C. Amazon EC2 Auto Scaling
D. AWS Batch
E. Amazon S3 website hosting
Answer: C,E (LEAVE A REPLY)

NEW QUESTION: 78
A company wants to move a web application to AWS. The application stores session information locally on each web server, which will make auto scaling difficult As part of the migration, the application will be rewritten to decouple the session data from the web servers. The company requires low latency, scalability, and availability. Which service will meet the requirements for storing the session information in the MOST cost-effective way?

A. Amazon ElastiCache with the Memcached engine
B. Amazon S3
C. Amazon RDS MySQL
D. Amazon ElastiCache with the Redis engine
Answer: (SHOW ANSWER)
Explanation
https://aws.amazon.com/caching/session-management/
Building real-time apps across versatile use cases like gaming, geospatial service, caching, session stores, or queuing, with advanced data structures, replication, and point-in-time snapshot support. Memcached: Building a simple, scalable caching layer for your data-intensive apps. https://aws.amazon.com/elasticache/

NEW QUESTION: 79
In AWS IAM, which of the following predefined policy condition keys checks how long ago (in seconds) the MFA-validated security credentials making the request were issued using multi-factor authentication (MFA)?
A. aws:MultiFactorAuthAge
B. aws:MultiFactorAuthLast
C. aws:MFAAge
D. aws:MultiFactorAuthPrevious
Answer: A (LEAVE A REPLY)
aws:MultiFactorAuthAge is one of the predefined keys provided by AWS that can be included within a Condition element of an IAM policy. The key allows to check how long ago (in seconds) the MFA-validated security credentials making the request were issued using Multi-Factor Authentication (MFA).
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

NEW QUESTION: 80
A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.
Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?
A. Create a transit VPC across two AZs using a third-party routing appliance. Create a VPN connection to each VPC. Default route internet traffic to the transit VPC.
B. Create multiple hosted-private AWS Direct Connect VIFs, one per account, each with a Direct Connect gateway. Default route internet traffic back to an on-premises router to route to the internet.
C. Create a central VPC for outbound internet traffic. Use VPC peering to default route to a set of redundant NAT gateway in the central VPC.

D. Create a proxy fleet in a central VPC account. Create an AWS PrivateLink endpoint service in the central VPC. Use PrivateLink interface for internet connectivity through the proxy fleet.

Answer: D (LEAVE A REPLY)

Explanation

user proxy fleet over PrivateLink. As explained in this AWS

website:https://aws.amazon.com/blogs/networking-and-content-delivery/how-to-use-aws-privatelink-to-secure-a

NEW QUESTION: 81

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: 82

An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.

Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs ?

A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.

B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.

C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.

D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.

Answer: A (LEAVE A REPLY)

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

NEW QUESTION: 83
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 that belong 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)

NEW QUESTION: 84
A company has Linux-based Amazon EC2 instances. Users must access the instances by using SSH with EC2 SSH Key pairs. Each machine requires a unique EC2 Key pair.
The company wants to implement a key rotation policy that will, upon request, automatically rotate all the EC2 key pairs and keep the key in a securely encrypted place. The company will accept less than 1 minute of downtime during key rotation.
Which solution will meet these requirement?
A. Store all the keys in AWS Secrets Manager. Define a Secrets Manager rotation schedule to invoke an AWS Lambda function to generate new key pairs. Replace public Keys on EC2 instances. Update the private keys in Secrets Manager.
B. Store all the keys in Parameter. Store, a capability of AWS Systems Manager, as a string. Define a Systems Manager maintenance window to invoke an AWS Lambda function to generate new key pairs.
Replace public keys on EC2 instance. Update the private keys in parameter.
C. Import the EC2 key pairs into AWS Key Management Service (AWS KMS). Configure automatic key rotation for these key pairs. Create an Amazon EventlBridge scheduled rule to invoke an AWS Lambda function to initiate the key rotation AWS KMS.

D. Add all the EC2 instances to Feet Manager, a capability of AWS Systems Manager. Define a Systems Manager maintenance window to issue a Systems Manager Run Command document to generate new Key pairs and to rotate public keys to all the instances in Feet Manager.

Answer: A (LEAVE A REPLY)

To meet the requirements for automatic key rotation of EC2 SSH key pairs with minimal downtime, storing the keys in AWS Secrets Manager and defining a rotation schedule is the most suitable solution. AWS Secrets Manager supports automatic rotation of secrets, including SSH keys, by invoking a Lambda function that can handle the creation of new key pairs and the replacement of public keys on EC2 instances. Updating the corresponding private keys in Secrets Manager ensures secure and centralized management of SSH keys, complying with the key rotation policy and minimizing operational overhead.

References:

AWS Secrets Manager Documentation: Describes how to store and rotate secrets, including SSH keys, using Secrets Manager and Lambda functions.

AWS Lambda Documentation: Provides information on creating Lambda functions for custom secret rotation logic.

AWS Best Practices for Security: Highlights the importance of key rotation and how AWS services like Secrets Manager can facilitate secure and automated key management.

NEW QUESTION: 85

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 J

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

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

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

E. Create an authorization lo associate the private hosted zone in Account A with the new VPC In Account B

Answer: C,E (LEAVE A REPLY)

NEW QUESTION: 86

A company wants to migrate its MySQL database from on premises to AWS. The company recently experienced a database outage that significantly impacted the business. To ensure this does not happen again, the company wants a reliable database solution on AWS that minimizes data loss and stores every transaction on at least two nodes.
Which solution meets these requirements?
A. Create an Amazon RDS MySQL DB instance and then create a read replica in a separate AWS Region that synchronously replicates the data.
B. Create an Amazon RDS DB instance with synchronous replication to three nodes in three Availability Zones.
C. Create an Amazon EC2 instance with a MySQL engine installed that triggers an AWS Lambda function to synchronously replicate the data to an Amazon RDS MySQL DB instance.
D. Create an Amazon RDS MySQL DB instance with Multi-AZ functionality enabled to synchronously replicate the data.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 87
A large company has a business-critical application that runs in a single AWS Region The application consists of multiple Amazon EC2 instances and an Amazon RDS Multi-AZ DB instance The EC2 instances run In an Amazon EC2 Auto Scaling group across multiple Availability Zones A solutions architect is implementing a disaster recovery (DR) plan for the application The solutions architect has created a pilot light application deployment in a new Region, which is referred to as the DR Region The DR environment has an Auto Scaling group with a single EC2 instance and a read replica of the RDS DB instance The solutions architect must automate a failover from the primary application environment to the pilot light environment in the DR Region Which solution meets these requirements with the MOST operational efficiency"
A. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region Configure the cron task to check whether the application is available Upon failure, the cron task modifies the DR environment by promoting the read replica and by adding EC2 instances to the Auto Scaling group
B. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the primary Region Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic in the DR Region Use an AWS Lambda function that is invoked by Amazon SNS in the DR Region to promote the read replica and to add EC2 instances to the Auto Scaling group
C. Create a cron task that runs every 5 minutes by using one of the application's EC2 instances in the primary Region Configure the cron task to check whether the application is available Upon failure, the cron task notifies a systems operator and attempts to restart the application services
D. Publish an application availability metric to Amazon CloudWatch in the DR Region from the application environment in the pnmary Region Create a CloudWatch alarm in the DR Region that is invoked when the application availability metric stops being delivered Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS> topic in the DR Region Add an email subscription to the SNS topic that sends messages to the application owner upon notification, instruct a systems operator to sign in to the AWS Management Console and initiate failover operations for the application

Answer: B (LEAVE A REPLY)

NEW QUESTION: 88

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: (SHOW ANSWER)

Explanation

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 durability . 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 to attach 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: 89
A company has created a VPC with multiple private subnets in multiple Availability Zones (AZs) and one public subnet in one of the AZs. The public subnet is used to launch a NAT gateway. There are instances in the private subnets that use a NAT gateway to connect to the internet. In case of an AZ failure, the company wants to ensure that the instances are not all experiencing internet connectivity issues and that there is a backup plan ready.
Which solution should a solutions architect recommend that is MOST highly available?
A. Create a new public subnet with a NAT gateway in the same AZ. Distribute the traffic between the two NAT gateways.

B. Create an Amazon EC2 NAT instance in the same public subnet. Replace the NAT gateway with the NAT instance and associate the instance with an Auto Scaling group with an appropriate scaling policy.
C. Create public subnets in each AZ and launch a NAT gateway in each subnet. Configure the traffic from the private subnets in each AZ to the respective NAT gateway.
D. Create an Amazon EC2 NAT instance in a new public subnet. Distribute the traffic between the NAT gateway and the NAT instance.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 90
A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:
* The data must be highly durable and available.
* The data must always be encrypted at rest and in transit.
* The encryption key must be managed by the company and rotated periodically.
Which of the following solutions should the Solutions Architect recommend?
A. Use Amazon DynamoDB with SSL to connect to DynamoDB. Use an AWS KMS key to encrypt DynamoDB objects at rest.
B. Deploy instances with Amazon EBS volumes attached to store this data. Use EBS volume encryption using an AWS KMS key to encrypt the data.
C. Deploy the storage gateway to AWS in file gateway mode. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
D. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
Answer: (SHOW ANSWER)

NEW QUESTION: 91
A company has a web service deployed in the following two AWS Regions: us-west-2 and us-est-
1. Each AWS region runs an identical version of the web service. Amazon Route 53 is used to route customers to the AWS Region that has the lowest latency.
The company wants to improve the availability of the web service in case an outage occurs in one of the two AWS Regions.
A Solutions Architect has recommended that a Route 53 health check be performed. The health check must detect a specific text on an endpoint.
What combination of conditions should the endpoint meet to pass the Route 53 health check?
(Choose two.)
A. The endpoint must respond to the request within the number of seconds specified when creating the health check.
B. The endpoint must establish a TCP connection within 10 seconds.
C. The endpoint must return an HTTP 200 status code.
D. The specific text string must appear within the first 5,120 bytes of the response.

E. The endpoint must return an HTTP 2xx or 3xx status code.
Answer: B,E (LEAVE A REPLY)

NEW QUESTION: 92
A company hosts its application in the AWS Cloud. The application runs on Amazon EC2 instances behind an Elastic Load Balancer in an Auto Scaling group and with an Amazon DynamoDB table. The company wants to ensure the application can be made available in another AWS Region with minimal downtime.
What should a solutions architect do to meet these requirements with the LEAST amount of downtime?
A. Create an AWS CloudFormation template to create EC2 instances and a load balancer to be executed when needed. Configure the DynamoDB table as a global table. Configure DNS failover to point to the new disaster recovery Region's load balancer.
B. Create an Auto Scaling group and load balancer in the disaster recovery Region. Configure the DynamoDB table as a global table. Create an Amazon CloudWatch alarm to trigger and AWS Lambda function that updates Amazon Route 53 pointing to the disaster recovery load balancer.
C. Create an AWS CloudFormation template to create EC2 instances, load balancers, and DynamoDB tables to be executed when needed. Configure DNS failover to point to the new disaster recovery Region's load balancer.
D. Create an Auto Scaling group and a load balancer in the disaster recovery Region. Configure the DynamoDB table as a global table. Configure DNS failover to point to the new disaster recovery Region's load balancer.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 93
How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?
A. Simply create a new volume in the other AZ and specify the original volume as the source.
B. Detach the volume and attach it to another EC2 instance in the other AZ.
C. Detach the volume, then use the ec2-migrate-voiume command to move it to another AZ.
D. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 94

A company hosts a blog post application on AWS using Amazon API Gateway. Amazon DynampDB, 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 are actively discussing topics in the comments section, and the company wants to increase use engagement by marking the comments appears in real time.
Which design should be used to reduce comment latency and improve user experience?
A. Use adge-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 WebSockts to deliver comments
D. Change the concurrency limit of the Lambda functions to lower the API response time.
Answer: (SHOW ANSWER)
Explanation
https://aws.amazon.com/appsync/

NEW QUESTION: 95
A company uses an Amazon EMR cluster to process data once a day. The raw data comes from Amazon S3, and the resulting processed data is also stored in Amazon S3. The processing must complete within 4 hours; currently, it only takes 3 hours. However, the processing time is taking 5 to 10 minutes. longer each week due to an increasing volume of raw data.
The team is also concerned about rising costs as the compute capacity increases. The EMR cluster is currently running on three m3 xlarge instances (one master and two core nodes).
Which of the following solutions will reduce costs related to the increasing compute needs?
A. Add additional task nodes, but use instance fleets with the master node in on-Demand mode and a mix of On-Demand and Spot Instances for the core and task nodes. Purchase a scheduled Reserved Instances for the master node.
B. Add additional task nodes, but use instance fleets with the master node in Spot mode and a mix of On-Demand and Spot Instances for the core and task nodes. Purchase enough scheduled Reserved Instances to offset the cost of running any On-Demand instances.
C. Add additional task nodes, but have the team purchase an all-upfront convertible Reserved Instance for each additional node to offset the costs.
D. Add additional task nodes, but use instance fleets with the master node in On-Demand mode and a mix of On-Demand and Spot Instances for the core and task nodes. Purchase a standard all- upfront Reserved Instance for the master node.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 96
A company is running several workloads in a single AWS account. A new company policy stales that engineers can provision only approved resources and that engineers must use AWS CloudFormation to provision these

resources. A solutions architect needs to create a solution to enforce the new restriction on the IAM role that the engineers use for access.
What should the solutions architect do to create the solution?
A. Provision resources in AWS CloudFormation stacks. Update the IAM policy for the engineers' IAM role to only allow access to their own AWS CloudFormation stack.
B. Update the IAM policy for the engineers" IAM role with permissions to only allow provisioning of approved resources and AWS CloudFormation. Use AWS CloudFormation templates to create stacks with approved resources.
C. Upload AWS CloudFormation templates that contain approved resources to an Amazon S3 bucket.
Update the IAM policy for the engineers' IAM role to only allow access to Amazon S3 and AWS CloudFormation.
Use AWS CloudFormation templates to provision resources.
D. Update the IAM policy for the engineers' IAM role with permissions to only allow AWS CloudFormation actions. Create a new IAM policy with permission to provision approved resources, and assign the policy to a new IAM service role. Assign the IAM service role to AWS CloudFormation during stack creation.
Answer: (SHOW ANSWER)

NEW QUESTION: 97
A company wants to allow its Marketing team to perform SQL queries on customer records to identify market segments. The data is spread across hundreds of files. The records must be encrypted in transit and at rest. The Team Manager must have the ability to manage users and groups, but no team members should have access to services or resources not required for the SQL queries. Additionally, Administrators need to audit the queries made and receive notifications when a query violates rules defined by the Security team.
AWS Organizations has been used to create a new account and an AWS IAM user with administrator permissions for the Team Manager.
Which design meets these requirements?
A. Apply a service control policy (SCP) that denies to all services except IAM, Amazon DynamoDB, and AWS CloudTrail. Store customer records in DynamoDB and train users to execute queries using the AWS CLI. Enable DynamoDB streams to track the queries that are issued and use an AWS Lambda function for real-time monitoring and alerting.
B. Apply a service control policy (SCP) that allows access to IAM, Amazon RDS, and AWS CloudTrail.
Load customer records in Amazon RDS MySQL and train users to execute queries using the AWS CLI. Stream the query logs to Amazon CloudWatch Logs from the RDS database instance. use a subscription filter with AWS lambda functions to audit and alarm on queries against personal data.
C. Apply a service control policy (SCP) that denies access to all services except IAM, Amazon Athena, Amazon S3, and AWS CloudTrail. Store customer record files in Amazon S3 and train users to execute queries using the CLI via Athena. Analyze CloudTrail events to audit and alarm on queries against personal data.
D. Apply a service control policy (SCP) that allows to IAM, Amazon Athena, Amazon S3, and AWS CloudTrail.
Store customer records as files in Amazon S3 and train users to leverage the Amazon S3 Select feature and execute queries using the AWS CLI. Enable S3 object-level logging and analyze CloudTrail events to audit and alarm on queries against personal data.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 98
A solutions architect has an operational workload deployed on Amazon EC2 instances in an Auto Scaling Group The VPC architecture spans two Availability Zones (AZ) with a subnet in each that the Auto Scaling group is targeting. The VPC is connected to an on-premises environment and connectivity cannot be interrupted The maximum size of the Auto Scaling group is 20 instances in service. The VPC IPv4 addressing is as follows:
VPCCIDR 10 0 0 0/23
AZ1 subnet CIDR: 10 0 0 0724
AZ2 subnet CIDR: 10.0.1 0724
Since deployment, a third AZ has become available in the Region The solutions architect wants to adopt the new AZ without adding additional IPv4 address space and without service downtime. Which solution will meet these requirements?
A. Update the Auto Scaling group to use the AZ2 subnet only Delete and re-create the AZ1 subnet using half the previous address space Adjust the Auto Scaling group to also use the new AZI subnet When the instances are healthy, adjust the Auto Scaling group to use the AZ1 subnet only Remove the current AZ2 subnet Create a new AZ2 subnet using the second half of the address space from the original AZ1 subnet Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets.
B. Terminate the EC2 instances in the AZ1 subnet Delete and re-create the AZ1 subnet using hall the address space. Update the Auto Scaling group to use this new subnet. Repeat this for the second AZ.
Define a new subnet in AZ3: then update the Auto Scaling group to target all three new subnets
C. Create a new VPC with the same IPv4 address space and define three subnets, with one for each AZ Update the existing Auto Scaling group to target the new subnets in the new VPC
D. Update the Auto Scaling group to use the AZ2 subnet only Update the AZ1 subnet to have halt the previous address space Adjust the Auto Scaling group to also use the AZ1 subnet again. When the instances are healthy, adjust the Auto Seating group to use the AZ1 subnet only. Update the current AZ2 subnet and assign the second half of the address space from the original AZ1 subnet Create a new AZ3 subnet using half the original AZ2 subnet address space, then update the Auto Scaling group to target all three new subnets
Answer: (SHOW ANSWER)
https://repost.aws/knowledge-center/vpc-ip-address-range

NEW QUESTION: 99
You have written a CloudFormation template that creates 1 Elastic Load Balancer fronting 2 EC2 Instances. Which section of the template should you edit so that the DNS of the load balancer is returned upon creation of the stack?
A. Parameters
B. Outputs
C. Mappings
D. Resources
Answer: (SHOW ANSWER)
Explanation

You can use AWS CloudFormation's sample templates or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application.
In the following example, the output named BackupLoadBalancerDNSName returns the DNS name for the resource with the logical ID BackupLoadBalancer only when the CreateProdResources condition is true. (The second output shows how to specify multiple outputs.) "Outputs" : {
"BackupLoadBalancerDNSName" : {
"Description": "The DNSName of the backup load balancer", "Value" : { "Fn::GetAtt" : [
"BackupLoadBalancer", "DNSName"]}, "Condition" : "CreateProdResources"
},
"InstanceID" : {
"Description": "The Instance ID", "Value" : { "Ref" : "EC2Instance" }
}
}
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/outputs-section-structure.html

NEW QUESTION: 100
A company is developing and hosting several projects in the AWS Cloud. The projects are developed across multiple AWS accounts under the same organization in AWS Organizations. The company requires the cost lor cloud infrastructure to be allocated to the owning project. The team responsible for all of the AWS accounts has discovered that several Amazon EC2 instances are lacking the Project tag used for cost allocation.
Which actions should a solutions architect take to resolve the problem and prevent it from happening in the future? (Select THREE.)
A. Create an AWS Config rule in each account to find resources with missing tags.
B. Create an SCP in the organization with a deny action for ec2:RunInstances if the Project tag is missing.
C. Use Amazon Inspector in the organization to find resources with missing tags.
D. Create an IAM policy in each account with a deny action for ec2:RunInstances if the Project tag is missing.
E. Create an AWS Config aggregator for the organization to collect a list of EC2 instances with the missing Project tag.
F. Use AWS Security Hub to aggregate a list of EC2 instances with the missing Project tag.
Answer: (SHOW ANSWER)
Explanation
https://docs.aws.amazon.com/config/latest/developerguide/config-rule-multi-account-deployment.html
https://docs.aws.amazon.com/config/latest/developerguide/aggregate-data.html
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_tagging.html

NEW QUESTION: 101
You have just added a new instance to your Auto Scaling group, which receives ELB health checks. An ELB heath check says the new instance's state is out of Service.
What does Auto Scaling do in this particular scenario?
A. It replaces the instance with a healthy one
B. It stops the instance

C. It marks an instance as unhealthy
D. It terminates the instance
Answer: (SHOW ANSWER)
Explanation
If you have attached a load balancer to your Auto Scaling group, you can have Auto Scaling include the results of Elastic Load Balancing health checks when it determines the health status of an instance. After you add ELB health checks, Auto Scaling will mark an instance as unhealthy if Elastic Load Balancing reports the instance state as Out of Service. Frequently, an Auto Scaling instance that has just come into service needs to warm up before it can pass the Auto Scaling health check. Auto Scaling waits until the health check grace period ends before checking the health status of the instance. While the EC2 status checks and ELB health checks can complete before the health check grace period expires, Auto Scaling does not act on them until the health check grace period expires. To provide ample warm-up time for your instances, ensure that the health check grace period covers the expected startup time for your application.
http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html

NEW QUESTION: 102
A company hosts its core network services, including directory services and DNS, in its on-premises data center. The data center is connected to the AWS Cloud using AWS Direct Connect (DX). Additional AWS accounts are planned that will require quick, cost-effective, and consistent access to these network services.
What should a solutions architect implement to meet these requirements with the LEAST amount of operational overhead?
A. Create a DX connection in each new account. Route the network traffic to the on-premises servers.
B. Configure AWS Transit Gateway between the accounts. Assign DX to the transit gateway and route network traffic to the on-premises servers.
C. Create a VPN connection between each new account and the DX VPC. Route the network traffic to the on-premises servers.
D. Configure VPC endpoints in the DX VPC for all required services. Route the network traffic to the on-premises servers.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 103
A company has a web application that securely uploads pictures and videos to an Amazon S3 bucket. The company requires that only authenticated users are allowed to post content. The application generates a presigned URL that is used to upload objects through a browser interface. Most users are reporting slow upload times for objects larger than 100 MB.
What can a Solutions Architect do to improve the performance of these uploads while ensuring only authenticated users are allowed to post content?
A. Set up an Amazon API Gateway with an edge-optimized API endpoint that has a resource as an S3 service proxy. Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using a COGNITO_USER_POOLS authorizer. Have the browser interface use API Gateway instead of the presigned URL to upload objects.

B. Set up an Amazon API Gateway with a regional API endpoint that has a resource as an S3 service proxy. Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using an AWS Lambda authorizer. Have the browser interface use API Gateway instead of the presigned URL to upload API objects.
C. Enable an S3 Transfer Acceleration endpoint on the S3 bucket. Use the endpoint when generating the presigned URL. Have the browser interface upload the objects to this URL using the S3 multipart upload API.
D. Configure an Amazon CloudFront distribution for the destination S3 bucket. Enable PUT and POST methods for the CloudFront cache behavior. Update the CloudFront origin to use an origin access identity (OAI). Give the OAI user s3:PutObject permissions in the bucket policy. Have the browser interface upload objects using the CloudFront distribution
Answer: (SHOW ANSWER)
Explanation: S3 Transfer Acceleration is a feature that enables fast, easy, and secure transfers of files over long distances between your client and an S3 bucket1. It works by leveraging the CloudFront edge network to route your requests to S3 over an optimized network path1. By using a Transfer Acceleration endpoint when generating a presigned URL, you can allow authenticated users to upload objects faster and more reliably2. Additionally, using the S3 multipart upload API can improve the performance of large object uploads by breaking them into smaller parts and uploading them in parallel3.
References:
S3 Transfer Acceleration
Using Transfer Acceleration with presigned URLs
Uploading objects using multipart upload API

NEW QUESTION: 104
Which EC2 functionality allows the user to place the Cluster Compute instances in clusters?
A. Cluster group
B. Cluster security group
C. GPU units
D. Cluster placement group
Answer: D (LEAVE A REPLY)
Explanation
The Amazon EC2 cluster placement group functionality allows users to group cluster compute instances in clusters.
https://aws.amazon.com/ec2/faqs/

NEW QUESTION: 105
A company is using AWS to run an internet-facing production application written in Node.js. The Development team is responsible for pushing new versions of their software directly to production. The application software is updated multiple times a day. The team needs guidance from a Solutions Architect to help them deploy the software to the production fleet quickly and with the least amount of disruption to the service.
Which option meets these requirements?

A. Use AWS Elastic Beanstalk to host the production application. For software changes, upload the new application version to Elastic Beanstalk to push this to the production fleet using a blue/green deployment method.
B. Use AWS CodeDeploy to push the prepackaged AMI to production. For software changes, reconfigure CodeDeploy with new AMI identification to push the new AMI to the production fleet.
C. Prepackage the software into an AMI and then use Auto Scaling to deploy the production fleet. For software changes, update the AMI and allow Auto Scaling to automatically push the new AMI to production.
D. Deploy the base AMI through Auto Scaling and bootstrap the software using user data. For software changes, SSH to each of the instances and replace the software with the new version.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 106
A company wants to launch an online shopping website in multiple countries and must ensure that customers are protected against potential "man-in-the-middle" attacks.
Which architecture will provide the MOST secure site access?
A. Use Route 53 for domain registration. Register 2048-bit encryption keys from a third-party certificate.
Use a third-party DNS service that supports DNSSEC for DNS requests that use the customer managed keys.
Import the customer managed keys to ACM to deploy the certificates to Classic Load Balancers configured with those TLS/SSL certificates for the site. Use the Server Name Identification extension in all clients requests to the site.
B. Use Route 53 for domain registration, and host the company DNS root servers on Amazon EC2 instances running Bind. Enable DNSSEC for DNS requests. Use ACM to register TLS/SSL certificates for the shopping website, and use Application Load Balancers configured with those TLS/SSL certificates for the site. Use the Server Name Identification extension in all client requests to the site.
C. Register 2048-bit encryption keys from a third-party certificate service. Use a third-party DNS provider that uses the customer managed keys for DNSSec. Upload the keys to ACM, and use ACM to automatically deploy the certificates for secure web services to an EC2 front-end web server fleet by using NGINX. Use the Server Name Identification extension in all client requests to the site.
D. Use Amazon Route 53 for domain registration and DNS services. Enable DNSSEC for all Route 53 requests. use AWS Certificate Manager (ACM) to register TLS/SSL certificates for the shopping website, and use Application Load Balancers configured with those TLS/SSL certificates for the site.
Use the Server Name Identification extension in all client requests to the site.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 107
A company has an Amazon EC2 deployment that has the following architecture:
* An application tier that contains 8 m4.xlarge instances
* A Classic Load Balancer
* Amazon S3 as a persistent data store
After one of the EC2 instances fails, users report very slow processing of their requests. A Solutions Architect must recommend design changes to maximize system reliability. The solution must minimize costs.
What should the Solution Architect recommend?
A. Migrate the existing EC2 instances to a serverless deployment using AWS Lambda functions
B. Change the Classic Load Balancer to an Application Load Balancer
C. Replace the application tier with m4.large instances in an Auto Scaling group
D. Replace the application tier with 4 m4.2xlarge instances
Answer: B (LEAVE A REPLY)
Explanation
By default, connection draining is enabled for Application Load Balancers but must be enabled for Classic Load Balancers. When Connection Draining is enabled and configured, the process of deregistering an instance from an Elastic Load Balancer gains an additional step. For the duration of the configured timeout, the load balancer will allow existing, in-flight requests made to an instance to complete, but it will not send any new requests to the instance. During this time, the API will report the status of the instance as InService, along with a message stating that "Instance deregistration currently in progress." Once the timeout is reached, any remaining connections will be forcibly closed.
https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.htmlhttps://aws.amazon.com/b

NEW QUESTION: 108
You are responsible for a legacy web application whose server environment is approaching end of life.
You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:
- the VM's single 10GB VMDK is almost full;
- the virtual network Interface still uses the 10Mbps dnver, which
leaves your 100Mbps WAN connection completely underutilized;
- it is currently running on a highly customized, Windows VM within a
VMware environment; ?you do not have the installation media.
This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours, RPO (Recovery Point Objective) of 1 hour.
How could you best migrate this application to AWS while meeting your business continuity requirements?
A. Use S3 to create a backup of the VM and restore the data into EC2.
B. Use the EC2 VM Import Connector for vCenter to import the VM into EC2.

C. Use the ec2-bundle-instance API to import an image of the VM into EC2.
D. Use Import/Export to import the VM as an EBS snapshot and attach to EC2.
Answer: B (LEAVE A REPLY)
https://aws.amazon.com/developertools/2759763385083070

NEW QUESTION: 109
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 that belong 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: C (LEAVE A REPLY)
Explanation
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

NEW QUESTION: 110
A company has an application that runs on Amazon EC2 instances. A solutions architect is designing VPC infrastructure in an AWS Region where the application needs to access an Amazon Aurora DB cluster. The EC2 instances are all associated with the same security group. The DB cluster is associated with its own security group.
The solutions architect needs to add rules to the security groups to provide the application with least privilege access to the DB cluster.
Which combination of steps will meet these requirements? (Select TWO.)

A. Add an inbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the source over the default Aurora port.
B. Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port.
C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port.
D. Add an outbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the destination over the default Aurora port.
E. Add an outbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the destination over the ephemeral ports.
Answer: A,B (LEAVE A REPLY)
B). Add an outbound rule to the EC2 instances' security group. Specify the DB cluster's security group as the destination over the default Aurora port. This allows the instances to make outbound connections to the DB cluster on the default Aurora port. C. Add an inbound rule to the DB cluster's security group. Specify the EC2 instances' security group as the source over the default Aurora port. This allows connections to the DB cluster from the EC2 instances on the default Aurora port.

NEW QUESTION: 111
An enterprise runs 103 line-of-business applications on virtual machines in an on-premises data center. Many of the applications are simple PHP, Java, or Ruby web applications, are no longer actively developed, and serve little traffic.
Which approach should be used to migrate these applications to AWS with the LOWEST infrastructure costs ?
A. Deploy the applications to single-instance AWS Elastic Beanstalk environments without a load balancer.
B. Use AWS SMS to create AMIs for each virtual machine and run them in Amazon EC2.
C. Convert each application to a Docker image and deploy to a small Amazon ECS cluster behind an Application Load Balancer.
D. Use VM Import/Export to create AMIs for each virtual machine and run them in single-instance AWS Elastic Beanstalk environments by configuring a custom image.
Answer: C (LEAVE A REPLY)
Explanation
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-types.html

NEW QUESTION: 112
A company needs to run a software package that has a license that must be run on the same physical host for the duration of Its use. The software package is only going to be used for 90 days The company requires patching and restarting of all instances every 30 days How can these requirements be met using AWS?
A. Run a dedicated instance with auto-placement disabled.
B. Run the instance on a dedicated host with Host Affinity set to Host.
C. Run an On-Demand Instance with a Reserved Instance to ensure consistent placement.
D. Run the instance on a licensed host with termination set for 90 days.

Answer: B (LEAVE A REPLY)

Explanation

Host Affinity is configured at the instance level. It establishes a launch relationship between an instance and a Dedicated Host. (This set which host the instance can run on) Auto-placement allows you to manage whether instances that you launch are launched onto a specific host, or onto any available host that has matching configurations. Auto-placement must be configured at the host level. (This sets which instance the host can run.) When affinity is set to Host, an instance launched onto a specific host always restarts on the same host if stopped. This applies to both targeted and untargeted launches.

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/how-dedicated-hosts-work.html When affinity is set to Off, and you stop and restart the instance, it can be restarted on any available host.

However, it tries to launch back onto the last Dedicated Host on which it ran (on a best-effort basis).

NEW QUESTION: 113

A finance company hosts a data lake in Amazon S3. The company receives financial data records over SFTP each night from several third parties. The company runs its own SFTP server on an Amazon EC2 instance in a public subnet of a VPC. After the files are uploaded, they are moved to the data lake by a cron job that runs on the same instance. The SFTP server is reachable on DNS sftp.example.com through the use of Amazon Route 53.

What should a solutions architect do to improve the reliability and scalability of the SFTP solution?

A. Migrate the SFTP server to a file gateway in AWS Storage Gateway. Update the DNS record sftp.example.com in Route 53 to point to the file gateway endpoint.

B. 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.

C. 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.

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: C (LEAVE A REPLY)

NEW QUESTION: 114

An elastic network interface (ENI) is a virtual network interface that you can attach to an instance in a VPC. An ENI can include one public IP address, which can be auto-assigned to the elastic network interface for eth0 when you launch an instance, but only when you_____.

A. create an elastic network interface for eth1

B. include a MAC address

C. use an existing network interface

D. create an elastic network interface for eth0

Answer: D (LEAVE A REPLY)

Explanation

An elastic network interface (ENI) is defined as a virtual network interface that you can attach to an instance in a VPC and can include one public IP address, which can be auto-assigned to the elastic network interface for eth0

when you launch an instance, but only when you create an elastic network interface for eth0 instead of using an existing network interface.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html

NEW QUESTION: 115
A company has an application that generates invoices and makes the invoices available online. Invoices are stored as PDFs in an Amazon S3 bucket. Customers typically only view each invoice during the month it is issued. However, past invoices need to be immediately available. There are concerns over rising storage costs as the company gains more customers.
What is the MOST cost-effective method to store the data?
A. Use Amazon S3 for current invoices. Set up lifecycle rules to migrate invoices to the GLACIER storage class after 30 days.
B. Store the invoices as text files. Use Amazon CloudFront to convert the invoices from text to PDF when customers download invoices.
C. Use Amazon S3 for current invoices. Set up lifecycle rules to migrate invoices to Amazon S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
D. Store the invoices as binaries in an Amazon RDS database instance. Retrieve them from the database when customers request invoices.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 116
An online magazine will launch its latest edition this month. This edition will be the first to be distributed globally. The magazine's dynamic website currently uses an Application Load Balance in front of the web tier, a fleet of Amazon EC2 instances for web and application servers, and Amazon Aurora MySQL. Portions of the website include static content and almost all traffic is read-only.
The magazine is exporting a significant spike in internet traffic when the new edition is launched. Optimal performance is a top priority for the week following the launch.
Which combination of steps should a solutions architect take to reduce system response times for a global audience? (Select Two.)
A. Use logical cross-Region replication to replicate the Aurora MySQL database to a secondary Region Replace the web servers with Amazon S3. Deploy S3 buckets in cross-Region replication mode.
B. Use an Aurora global database for physical cross-Region replication. Use Amazon S3 with cross-Region replication tor static content and resources. Deploy the web and application tiers in Regions across the world.
C. Migrate the database from Amazon Aurora to Amazon RDS tor MySQL Ensure all three of the application tiers-web. application, and database-are in private subnets.
D. Introduce Amazon Route 53 with latency-based routing and Amazon CloudFront distributions. Ensure the web and application tiers are each in Auto Scaling groups.
E. Ensure the web and application tiers are each in Auto Scaling groups. Introduce an AWS Direct Connect connection Deploy the web and application liars in Regions across the world.
Answer: (SHOW ANSWER)

NEW QUESTION: 117
In Amazon CloudWatch, you can publish your own metrics with the put-metric-data command.
When you create a new metric using the put-metric-data command, it can take up to two minutes before you can retrieve statistics on the new metric using the get-metric-statistics command. How long does it take before the new metric appears in the list of metrics retrieved using the list- metrics command?
A. After 2 minutes
B. Up to 15 minutes
C. More than an hour
D. Within a minute
Answer: B (LEAVE A REPLY)
You can publish your own metrics to CloudWatch with the put-metric-data command (or its Query API equivalent PutMetricData). When you create a new metric using the put-metric-data command, it can take up to two minutes before you can retrieve statistics on the new metric using the get-metric-statistics command. However, it can take up to fifteen minutes before the new metric appears in the list of metrics retrieved using the list-metrics command.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/publishingMetrics.html

NEW QUESTION: 118
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 (NIB) foe scalability.
Up to 10 business unit VPCs will need to be connected to the shared VPC Some of 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 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
B. 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
C. 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
D. 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 service console
Answer: C (LEAVE A REPLY)

NEW QUESTION: 119
A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.
A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.
Which solution will provide the LARGEST overall cost reduction while meeting these requirements?
A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.
B. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template. Use the EBS volume as the shared storage for the duration of the job. Detach the EBS volume when the job is complete.
C. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.
D. Migrate the data from the existing shared file system to an Amazon S3 bucket. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job. Delete the file gateway when the job is complete.
Answer: (SHOW ANSWER)
https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and-a

NEW QUESTION: 120
A company has AWS accounts that are in an organization in AWS rganizations. 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 amount to be 10% more than the reported average usage for the last 30 days from AWS Cost Explorer.
B. Configure an alert to notify the architecture team if the usage threshold is met. 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 EC2 usage 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 if Detective identifies a usage anomaly of more than 10%.
Answer: B (LEAVE A REPLY)
The correct answer is B.
B). This solution meets the requirements because it uses AWS Cost Anomaly Detection, which is a feature of AWS Cost Management that uses machine learning to identify and alert on anomalous spend and usage patterns. By configuring a monitor type of AWS Service and applying a filter of Amazon EC2, the solution can track the EC2 usage as a metric across the organization's accounts. By configuring an alert subscription with a threshold of 10%, the solution can notify the architecture team via email or Amazon SNS if the EC2 usage is more than 10% higher than the average usage for the last 30 days12
A). This solution is incorrect because it uses AWS Budgets, which is a feature of AWS Cost Management that helps to plan and track costs and usage. However, AWS Budgets does not support usage type of EC2 running hours as a budget type. The only supported usage types are Amazon S3 storage, Amazon EC2 RI utilization, and Amazon EC2 RI coverage. Moreover, AWS Budgets does not support setting the budget amount based on the reported average usage from AWS Cost Explorer. The budget amount has to be a fixed or variable value34
C). This solution is incorrect because it uses AWS Trusted Advisor, which is a feature of AWS Premium Support that provides recommendations to follow best practices for cost optimization, security, performance, and fault tolerance. However, AWS Trusted Advisor does not support configuring custom alerts based on EC2 usage or average usage for the last 30 days. The only supported alerts are based on predefined checks and thresholds that are applied to all services and resources in the account56
D). This solution is incorrect because it uses Amazon Detective, which is a service that helps to analyze and visualize security data to investigate potential security issues. However, Amazon Detective does not support configuring EC2 usage anomaly alerts based on average usage for the last 30 days. The only supported alerts are based on GuardDuty findings and other security-related events that are detected by machine learning models78 References:
1: AWS Cost Anomaly Detection - Amazon Web Services 2: Getting started with AWS Cost Anomaly Detection 3: Set Custom Cost and Usage Budgets - AWS Budgets - Amazon Web Services 4: Creating a budget - AWS Cost Management 5: AWS Trusted Advisor 6: AWS Trusted Advisor - AWS Support 7:
Security Investigation Visualization - Amazon Detective - AWS 8: What is Amazon Detective? - Amazon Detective

NEW QUESTION: 121
A company is serving files to its customers through an SFTP server that is accessible over the internet The SFTP server is running on a single Amazon EC2 instance with an Elastic IP address attached Customers connect to the SFTP server through its Elastic IP address and use SSH for authentication The EC2 instance also has an attached security group that allows access from all customer IP addresses.

A solutions architect must implement a solution to improve availability minimize the complexity of infrastructure management and minimize the disruption to customers who access files. The solution must not change the way customers connect Which solution will meet these requirements?

A. Disassociate the Elastic IP address from the EC2 instance Create an Amazon S3 bucket to be used for SFTP file hosting Create an AWS Transfer Family server. Configure the Transfer Family server with a publicly accessible endpoint Associate the SFTP Elastic IP address with the new endpoint. Point the Transfer Family server to the S3 bucket Sync all files from the SFTP server to the S3 bucket.

B. Disassociate the Elastic IP address from the EC2 instance Create an Amazon S3 bucket to be used for SFTP file hosting Create an AWS Transfer Family Server Configure the Transfer Family server with a VPC-hosted, internet-facing endpoint Associate the SFTP Elastic IP address with the new endpoint Attach the security group with customer IP addresses to the new endpoint Point the Transfer Family server to the S3 bucket. Sync all files from the SFTP server to the S3 bucket.

C. Disassociate the Elastic IP address from the EC2 instance. Create a new Amazon Elastic File System (Amazon EFS) file system to be used for SFTP file hosting. Create an AWS Fargate task definition to run an SFTP server Specify the EFS file system as a mount in the task definition Create a Fargate service by using the task definition, and place a Network Load Balancer (NLB) in front of the service.

When configuring the service, attach the security group with customer IP addresses to the tasks that run the SFTP server Associate the Elastic IP address with the NLB Sync all files from the SFTP server to the S3 bucket.

D. Disassociate the Elastic IP address from the EC2 instance. Create a multi-attach Amazon Elastic Block Store (Amazon EBS) volume to be used for SFTP file hosting. Create a Network Load Balancer (NLB) with the Elastic IP address attached. Create an Auto Scaling group with EC2 instances that run an SFTP server. Define in the Auto Scaling group that instances that are launched should attach the new multi-attach EBS volume Configure the Auto Scaling group to automatically add instances behind the NLB. configure the Auto Scaling group to use the security group that allows customer IP addresses for the EC2 instances that the Auto Scaling group launches Sync all files from the SFTP server to the new multi-attach EBS volume.

Answer: B (LEAVE A REPLY)

https://aws.amazon.com/premiumsupport/knowledge-center/aws-sftp-endpoint-type/

NEW QUESTION: 122

A user is trying to understand the detailed CloudWatch monitoring concept. Which of the below mentioned services does not provide detailed monitoring with CloudWatch?

A. AWS RDS
B. AWS ELB
C. AWS Route53
D. AWS EMR
Answer: D (LEAVE A REPLY)
CloudWatch is used to monitor AWS as well as the custom services. It provides either basic or detailed monitoring for the supported AWS products. In basic monitoring, a service sends data points to CloudWatch every five minutes, while in detailed monitoring a service sends data points to CloudWatch every minute. Services, such as RDS, EC2, Auto Scaling, ELB, and Route 53 can provide the monitoring data every minute.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/supported_services.htm l

NEW QUESTION: 123
A solutions architect is deploying a distributed database on multiple Amazon EC2 instances. The database stores all data on multiple instances so it can withstand the loss of an instance. The database requires block storage with latency and throughput to support several million transactions per second per server.
Which storage solution should the solutions architect use?
A. Amazon EBS
B. Amazon EC2 instance store
C. Amazon EFS
D. Amazon S3
Answer: (SHOW ANSWER)
Explanation/Reference: https://aws.amazon.com/ebs/

NEW QUESTION: 124
A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The quality assurance (QA) department needs to launch a large number of short-lived environments to test the application.
The application environments are currently launched by the manager of the department using an AWS CloudFormation template To launch the stack, the manager uses a role with permission to use CloudFormation EC2. and Auto Scaling APIs. The manager wants to allow testers to launch their own environments, but does not want to grant broad permissions to each user Which set up would achieve these goals?
A. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to assume the manager's role and add a policy that restricts the permissions to the template and the resources it creates Train users to launch the template from the CloudFormation console
B. Create an AWS Service Catalog product from the environment template Add a launch constraint to the product with the existing role Give users in the QA department permission to use AWS Service Catalog APIs only_ Train users to launch the template from the AWS Service Catalog console.
C. Upload the AWS CloudFormation template to Amazon S3 Give users in the QA department permission to use CloudFormation and S3 APIs, with conditions that restrict the permissions to the template and the resources it creates Train users to launch the template from the CloudFormation console.

D. Create an AWS Elastic Beanstalk application from the environment template Give users in the QA department permission to use Elastic Beanstalk permissions only Train users to launch Elastic Beanstalk environments with the Elastic Beanstalk CLI, passing the existing role to the environment as a service role

Answer: B (LEAVE A REPLY)

Explanation

https://docs.aws.amazon.com/iot/latest/developerguide/iot-rules.html

NEW QUESTION: 125

An ecommerce website is deploying its web application as Amazon Elastic Container Service (Amazon ECS) container instances behind an Application Load Balancer (ALB). During periods of high activity, the website slows down and availability is reduced. A solutions architect uses Amazon CloudWatch alarms to receive notifications whenever there is an availability issue so they can scale out resources. Company management wants a solution that automatically responds to such events.

Which solution meets these requirements?

A. Set up AWS Auto Scaling to scale out the ECS service when the service's CPU utilization is too high. Set up AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.

B. Set up AWS Auto Scaling to scale out the ECS service when the ALB CPU utilization is too high. Setup AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.

C. Set up AWS Auto Scaling to scale out the ECS service when the ALB target group CPU utilization is too high. Set up AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.

D. Set up AWS Auto Scaling to scale out the ECS service when there are timeouts on the ALB. Set up AWS Auto Scaling to scale out the ECS cluster when the CPU or memory reservation is too high.

Answer: (SHOW ANSWER)

NEW QUESTION: 126

In CloudFormation, if you want to map an Amazon Elastic Block Store to an Amazon EC2 instance, _____.

A. you reference the logical IDs to associate the block stores with the instance

B. you reference the physical IDs of the instance along with the resource type

C. you reference the instance IDs of the block store along with the resource properties

D. you reference the physical IDs of both the block stores and the instance

Answer: A (LEAVE A REPLY)

Explanation

In AWS CloudFormation, if you want to map an Amazon Elastic Block Store to an Amazon EC2 instance, you reference the logical IDs to associate the block stores with the instance.

http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/concept-resources.html

NEW QUESTION: 127

A company runs a Windows Server host in a public subnet that is configured to allow a team of administrators to connect over RDP to troubleshoot issues with hosts in a private subnet. The host must be available at all times

outside of a scheduled maintenance window, and needs to receive the latest operating system updates within 3 days of release.
What should be done to manage the host with the LEAST amount of administrative effort?
A. Run the host in a single-instance AWS Elastic Beanstalk environment. Configure the environment with a custom AMI to use a hardened machine image from AWS Marketplace. Apply system updates with AWS Systems Manager Patch Manager.
B. Run the host on AWS WorkSpaces. Use Amazon WorkSpaces Application Manager (WAM) to harden the host. Configure Windows automatic updates to occur every 3 days.
C. Run the host in an Auto Scaling group with a minimum and maximum instance count of 1. Use a hardened machine image from AWS Marketplace. Apply system updates with AWS Systems Manager Patch Manager.
D. Run the host in AWS OpsWorks Stacks. Use a Chief recipe to harden the AMI during instance launch. Use an AWS Lambda scheduled event to run the Upgrade Operating System stack command to apply system updates.
Answer: A (LEAVE A REPLY)
A: uses beanstalk env. creating a custom ami with option to use ssm.
B. Workspaces do no support windows server
C. there is nothing called hardened machine image in marketplace
D. opsworks command supports only linux

NEW QUESTION: 128
A company that hosts its web application on AWS wants to ensure all Amazon EC2 instances, Amazon RDS DB instances, and Amazon Redshift clusters are configured with tags. The company wants to minimize the effort of configuring and operating this check.
What should a solutions architect do to accomplish this?
A. Use AWS Config rules to define and detect resources that are not properly tagged.
B. Use Cost Explorer to display resources that are not properly tagged. Tag those resources manually.
C. Write API calls to check all resources for proper tag allocation. Periodically run the code on an EC2 instance.
D. Write API calls to check all resources for proper tag allocation. Schedule an AWS Lambda function through Amazon CloudWatch to periodically run the code.
Answer: (SHOW ANSWER)
Explanation/Reference: https://d1.awsstatic.com/whitepapers/aws-tagging-best-practices.pdf

NEW QUESTION: 129
A Company has a security event whereby an Amazon S3 bucket with sensitive information was made public. Company policy is to never have public S3 objects, and the Compliance team must be informed immediately when any public objects are identified.
How can the presence of a public S3 object be detected, set to trigger alarm notifications, and automatically remediated in the future? (Choose two.)
A. Turn on object-level logging for Amazon S3. Turn on Amazon S3 event notifications to notify by using an Amazon SNS topic when a PutObject API call is made with a public-read permission.

B. Configure an Amazon CloudWatch Events rule that invokes an AWS Lambda function to secure the S3 bucket.
C. Use the S3 bucket permissions for AWS Trusted Advisor and configure a CloudWatch event to notify by using Amazon SNS.
D. Turn on object-level logging for Amazon S3. Configure a CloudWatch event to notify by using an SNS topic when a PutObject API call with public-read permission is detected in the AWS CloudTrail logs.
E. Schedule a recursive Lambda function to regularly change all object permissions inside the S3 bucket.
Answer: B,D (LEAVE A REPLY)
https://aws.amazon.com/blogs/security/how-to-detect-and-automatically-remediate-unintended-permissions-in-amazon-s3-object-acls-with-cloudwatch-events/

NEW QUESTION: 130
A company requires that all internal application connectivity use private IP addresses. To facilitate this policy, a solutions architect has created interface endpoints to connect to AWS public services. Upon testing, the solutions architect notices that the service names are resolving to public IP addresses, and that internal services cannot connect to the interface endpoints.
Which step should the solutions architect take to resolve this issue?
A. Enable the private DNS option on the VPC attributes.
B. Configure the security group on the interface endpoint to allow connectivity to the AWS services.
C. Update the subnet route table with a route to the interface endpoint.
D. Configure an Amazon Route 53 private hosted zone with a conditional forwarder for the internal application.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 131
A user is trying to create a vault in AWS Glacier. The user wants to enable notifications.
In which of the below mentioned options can the user enable the notifications from the AWS console?
A. Glacier does not support the AWS console
B. Archival Upload Complete
C. Vault Upload Job Complete
D. Vault Inventory Retrieval Job Complete
Answer: D (LEAVE A REPLY)
Explanation
From AWS console the user can configure to have notifications sent to Amazon Simple Notifications Service (SNS). The user can select specific jobs that, on completion, will trigger the notifications such as Vault Inventory Retrieval Job Complete and Archive Retrieval Job Complete.
http://docs.aws.amazon.com/amazonglacier/latest/dev/configuring-notifications-console.html

NEW QUESTION: 132
A media company has a 30-TB repository of digital news videos. These videos are stored on tape in an on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an automated fashion and put them into a searchable catalog by using a

MAM feature. The company must be able to search based on information in the video, such as objects, scenery items, or people's faces. A catalog is available that contains faces of people who have appeared in the videos that include an image of each person. The company would like to migrate these videos to AWS.
The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system.
How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system?
A. Set up an AWS Storage Gateway, file gateway appliance on-premises. Use the MAM solution to extract the videos from the current archive and push them into the file gateway. Use the catalog of faces to build a collection in Amazon Rekognition. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata, and push the metadata into the MAM solution.
B. Set up an AWS Storage Gateway, tape gateway appliance on-premises. Use the MAM solution to extract the videos from the current archive and push them into the tape gateway. Use the catalog of faces to build a collection in Amazon Rekognition. Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Amazon Rekognition process the video in the tape gateway, retrieve the required metadata, and push the metadata into the MAM solution.
C. Configure a video ingestion stream by using Amazon Kinesis Video Streams. Use the catalog of faces to build a collection in Amazon Rekognition. Stream the videos from the MAM solution into Kinesis Video Streams. Configure Amazon Rekognition to process the streamed videos. Then, use a stream consumer to retrieve the required metadata, and push the metadata into the MAM solution. Configure the stream to store the videos in Amazon S3.
D. Set up an Amazon EC2 instance that runs the OpenCV libraries. Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volume mounted on this EC2 instance. Process the videos to retrieve the required metadata, and push the metadata into the MAM solution while also copying the video files to an Amazon S3 bucket.
Answer: C (LEAVE A REPLY)
https://docs.aws.amazon.com/rekognition/latest/dg/streaming-video.html

NEW QUESTION: 133
A solutions architect is building a web application that uses an Amazon RDS for PostgreSQL DB instance The DB instance is expected to receive many more reads than writes The solutions architect needs to ensure that the large amount of read traffic can be accommodated and that the DB instance is highly available.
Which steps should the solutions architect take to meet these requirements? (Select THREE.)
A. Create multiple read replicas and put them into an Auto Scaling group
B. Create multiple read replicas in different Availability Zones.
C. Create an Amazon Route 53 hosted zone and a record set for each read replica with a TTL and a weighted routing policy
D. Create an Application Load Balancer (ALBJ and put the read replicas behind the ALB.
E. Configure an Amazon CloudWatch alarm to detect a failed read replica Set the alarm to directly invoke an AWS Lambda function to delete its Route 53 record set.

F. Configure an Amazon Route 53 health check for each read replica using its endpoint

Answer: (SHOW ANSWER)

Explanation

https://aws.amazon.com/premiumsupport/knowledge-center/requests-rds-read-replicas/ You can use Amazon Route 53 weighted record sets to distribute requests across your read replicas. Within a Route 53 hosted zone, create individual record sets for each DNS endpoint associated with your read replicas and give them the same weight. Then, direct requests to the endpoint of the record set. You can incorporate Route 53 health checks to be sure that Route 53 directs traffic away from unavailable read replicas

NEW QUESTION: 134

A company that is new to AWS reports it has exhausted its service limits across several accounts that are on the Basic Support plan. The company would like to prevent this from happening in the future.

What is the MOST efficient way of monitoring and managing all service limits in the company's accounts?

A. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, provide notifications using Amazon SNS if the limits are close to exceeding the threshold.

B. Reach out to AWS Support to proactively increase the limits across all accounts. That way, the customer avoids creating and managing infrastructure just to raise the service limits.

C. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, programmatically increase the limits that are close to exceeding the threshold.

D. Use Amazon CloudWatch and AWS Lambda to periodically calculate the limits across all linked accounts using AWS Trusted Advisor, and use Amazon SNS for notifications if a limit is close to exceeding the threshold. Ensure that the accounts are using the AWS Business Support plan at a minimum.

Answer: D (LEAVE A REPLY)

Explanation

https://github.com/awslabs/aws-limit-monitor

https://aws.amazon.com/solutions/limit-monitor/

NEW QUESTION: 135

An organization has hosted an application on the EC2 instances. There will be multiple users connecting to the instance for setup and configuration of application. The organization is planning to implement certain security best practices. Which of the below mentioned pointers will not help the organization achieve better security arrangement?

A. Allow only IAM users to connect with the EC2 instances with their own secret access key.

B. Create a procedure to revoke the access rights of the individual user when they are not required to connect to EC2 instance anymore for the purpose of application configuration.

C. Apply the latest patch of OS and always keep it updated.

D. Disable the password based login for all the users. All the users should use their own keys to connect with the instance securely.

Answer: A (LEAVE A REPLY)

Since AWS is a public cloud any application hosted on EC2 is prone to hacker attacks. It becomes extremely important for a user to setup a proper security mechanism on the EC2 instances. A few of the security measures are listed below:
Always keep the OS updated with the latest patch
Always create separate users with in OS if they need to connect with the EC2 instances, create their keys and disable their password Create a procedure using which the admin can revoke the access of the user when the business work on the EC2 instance is completed Lock down unnecessary ports Audit any proprietary applications that the user may be running on the EC2 instance Provide temporary escalated privileges, such as sudo for users who need to perform occasional privileged tasks The IAM is useful when users are required to work with AWS resources and actions, such as launching an instance. It is not useful to connect (RDP / SSH) with an instance.
http://aws.amazon.com/articles/1233/

NEW QUESTION: 136
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: B (LEAVE A REPLY)
https://aws.amazon.com/about-aws/whats-new/2017/11/aws-lambda-supports-traffic-shifting-and- phased-deployments-with-aws-codedeploy/

NEW QUESTION: 137
A company recently deployed an application on AWS. The application uses Amazon DynamoDB. The company measured the application load and configured the RCUs and WCUs on the DynamoDB table to match the expected peak load. The peak load occurs once a week for a 4-hour period and is double the average load. The application load is close to the average load 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: 138
A company has an Amazon EC2 deployment that has the following architecture:
An application tier that contains 8 m4.xlarge instances
A Classic Load Balancer
Amazon S3 as a persistent data store
After one of the EC2 instances fails, users report very slow processing of their requests. A Solutions Architect must recommend design changes to maximize system reliability. The solution must minimize costs.
What should the Solution Architect recommend?
A. Migrate the existing EC2 instances to a serverless deployment using AWS Lambda functions
B. Change the Classic Load Balancer to an Application Load Balancer
C. Replace the application tier with m4.large instances in an Auto Scaling group
D. Replace the application tier with 4 m4.2xlarge instances
Answer: B (LEAVE A REPLY)
Explanation

By default, connection draining is enabled for Application Load Balancers but must be enabled for Classic Load Balancers. When Connection Draining is enabled and configured, the process of deregistering an instance from an Elastic Load Balancer gains an additional step. For the duration of the configured timeout, the load balancer will allow existing, in-flight requests made to an instance to complete, but it will not send any new requests to the instance. During this time, the API will report the status of the instance as InService, along with a message stating that "Instance deregistration currently in progress." Once the timeout is reached, any remaining connections will be forcibly closed.
https://docs.aws.amazon.com/autoscaling/ec2/userguide/attach-load-balancer-asg.html
https://aws.amazon.com/blogs/aws/elb-connection-draining-remove-instances-from-service-with-care/

NEW QUESTION: 139
One of the AWS account owners faced a major challenge in June as his account was hacked and the hacker deleted all the data from his AWS account. This resulted in a major blow to the business.
Which of the below mentioned steps would not have helped in preventing this action?
A. Setup an MFA for each user as well as for the root account user.
B. Take a backup of the critical data to offsite / on premise.
C. Create an AMI and a snapshot of the data at regular intervals as well as keep a copy to separate regions.
D. Do not share the AWS access and secret access keys with others as well do not store it inside programs, instead use IAM roles.
Answer: C (LEAVE A REPLY)
AWS security follows the shared security model where the user is as much responsible as Amazon. If the user wants to have secure access to AWS while hosting applications on EC2, the first security rule to follow is to enable MFA for all users. This will add an added security layer. In the second step, the user should never give his access or secret access keys to anyone as well as store inside programs. The better solution is to use IAM roles. For critical data of the organization, the user should keep an offsite/ in premise backup which will help to recover critical data in case of security breach.
It is recommended to have AWS AMIs and snapshots as well as keep them at other regions so that they will help in the DR scenario. However, in case of a data security breach of the account they may not be very helpful as hacker can delete that.
Therefore ,creating an AMI and a snapshot of the data at regular intervals as well as keep a copy to separate regions, would not have helped in preventing this action.
Reference: http://media.amazonwebservices.com/pdf/AWS_Security_Whitepaper.pdf

NEW QUESTION: 140
A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes. Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.

Which option meets the requirements with the LEAST disruption?
A. Create an AWS account for each business unit. Move each business unit's instances to its own account and set up a federation to allow users to access their business unit's account.
B. Set up a federation to allow users to use their corporate credentials, and lock the users down to their own VPC. Use a network ACL to block each VPC from accessing other VPCs.
C. Implement a tagging policy based on business units. Create an IAM policy so that each user can terminate instances belonging to their own business units only.
D. Set up role-based access for each user and provide limited permissions based on individual roles and the services for which each user is responsible.
Answer: C (LEAVE A REPLY)
Explanation
Principal - Control what the person making the request (the principal) is allowed to do based on the tags that are attached to that person's IAM user or role. To do this, use the aws:PrincipalTag/key-name condition key to specify what tags must be attached to the IAM user or role before the request is allowed.https://docs.aws.amazon.com/IAM/latest/UserGuide/access_iam-tags.html

NEW QUESTION: 141
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: C (LEAVE A REPLY)
Explanation
This solution would achieve the company's goal with the least operational overhead because it uses the AWS DMS service to replicate the data from the on-premises MySQL database to the target Aurora MySQL DB cluster

in the AWS cloud. The DMS service also provides a way to keep the data synchronized with change data capture (CDC) task. Additionally, the data migration process is simplified with the use of the AWS SCT.
Once the data is replicated, the rest of the application can be installed on EC2 instances by starting with a compatible base Amazon Machine Image (AMI). This eliminates the need for complex replication setup or regular failover and fallback drills.
https://aws.amazon.com/dms/

NEW QUESTION: 142
A company is developing a new on-demand video application that is based on microservices. The application will have 5 million users at launch and will have 30 million users after 6 months. The company has deployed the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. The company developed the application by using ECS services that use the HTTPS protocol.
A solutions architect needs to implement updates to the application by using blue/green deployments. The solution must distribute traffic to each ECS service through a load balancer. The application must automatically adjust the number of tasks in response to an Amazon CloudWatch alarm.
Which solution will meet these requirements?
A. Configure the ECS services to use the blue/green deployment type and a Network Load Balancer. Request increases to the service quota for tasks per service to meet the demand.
B. Configure the ECS services to use the blue/green deployment type and a Network Load Balancer. Implement an Auto Scaling group for each ECS service by using the Cluster Autoscaler.
C. Configure the ECS services to use the blue/green deployment type and an Application Load Balancer. Implement an Auto Seating group for each ECS service by using the Cluster Autoscaler.
D. Configure the ECS services to use the blue/green deployment type and an Application Load Balancer. Implement Service Auto Scaling for each ECS service.
Answer: (SHOW ANSWER)
https://repost.aws/knowledge-center/ecs-fargate-service-auto-scaling

NEW QUESTION: 143
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/Jpostld) to get post details
GET /users/{userld}. to get user details
GET /comments/{commentld}: 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 making the comments appear in real time Which design should be used to reduce comment latency and improve user experience?
A. Use AWS AppSync and leverage WebSockets to deliver comments
B. Change the concurrency limit of the Lambda functions to lower the API response time.
C. Modify the blog application code to request GET/commentsV{commentld} every 10 seconds
D. Use edge-optimized API with Amazon CloudFront to cache API responses.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 144
A retail company has structured its AWS accounts to be part of an organization in AWS Organizations. The company has set up consolidated billing and has mapped its departments to the following OUs: Finance, Sales, Human Resources (HR), Marketing, and Operations. Each OU has multiple AWS accounts, one for each environment within a department. These environments are development, test, pre-production, and production. The HR department is releasing a new system that will launch in 3 months. In preparation, the HR department has purchased several Reserved Instances (RIs) in its production AWS account. The HR department will install the new application on this account. The HR department wants to make sure that other departments cannot share the RI discounts.
Which solution will meet these requirements?
A. In the AWS Billing and Cost Management console for the HR department's production account turn off RI sharing.
B. Remove the HR department's production AWS account from the organization. Add the account 10 the consolidating billing configuration only.
C. In the AWS Billing and Cost Management console, use the organization's management account 10 turn off RI Sharing for the HR departments production AWS account.
D. Create an SCP in the organization to restrict access to the RIs. Apply the SCP to the OUs of the other departments.
Answer: C (LEAVE A REPLY)
Explanation
This solution will meet the HR department's requirements because it allows the company to turn off RI sharing for the specific production account of the HR department through the organization's management account.
This will prevent the reserved instances from being shared across the organization, and other departments will not be able to access or use the RIs. This approach allows the organization to manage the sharing of RIs at a central level, which could be useful if there are multiple departments or accounts that need to be managed in this way.

NEW QUESTION: 145
Which of the following cannot be used to manage Amazon ElastiCache and perform administrative tasks?
A. AWS software development kits (SDKs)
B. Amazon S3
C. ElastiCache command line interface (CLI)
D. AWS CloudWatch
Answer: D (LEAVE A REPLY)
CloudWatch is a monitoring tool and doesn't give users access to manage Amazon ElastiCache.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.Managing.html

NEW QUESTION: 146
A financial services company logs personality identifiable information to its application logs stored in Amazon S3. Due to regulatory compliance requirements, the log files must be encrypted at rest. The Security team has

mandated that the company's on-premises hardware security modules (HSMs) be used to generate the CMK material.
Which steps should the Solution Architected take to meet these requirements?
A. Provision AN AWS Direct Connect connection, ensuring there is no overlap of the RFC 1918 address space between on-premises hardware and the VPC. Configure an AWS bucket policy on the logging bucket requires all objects to be key material, and create a unique CMK for each logging event.
B. Create a CMK in AWS KMS with no key material and an origin of EXTERNAL. Import the key material generated from the on-premises HSMs into the CMK using the public key and import token provided by AWS. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
C. Create a new CMK in AWS KMS with AWS-provided key material and an origin of AWS-KMS.
Disable this CMK, and overwrite the key material with the material from the on-premises HSM using the public key and import token provided by AWS Re-enable the CMK. Enable automatic, key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the logging bucket that disallows uploads of non-encrypted data and requires that the encryption source be AWS KMS.
D. Create an AWS CloudHSM cluster. Create a new CMK in AWS KMS using AWS_CloudHSM as the source for the key material and an origin of AWS-CLOUDHSM. Enable automatic key rotation on the CMK with a duration of 1 year. Configure a bucket policy on the logging bucket the disallow uploads of unencrypted data and requires that the encryption source be AWS KMS.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 147
A company wants to migrate its web application to AWS. The legacy web application consists of a web tier, an application tier, and a MySQL database. The re-architected application must consist of technologies that do not require the administration team to manage instances or clusters.
Which combination of services should a solutions architect include in the overall architecture? (Choose two.)
A. Amazon Aurora Serverless
B. Amazon EC2 Spot Instances
C. Amazon Elasticsearch Service (Amazon ES)
D. AWS Fargate
E. Amazon RDS for MySQL
Answer: (SHOW ANSWER)

NEW QUESTION: 148
A company is running an application that uses an Amazon ElastiCache for Redis cluster as a caching layer A recent security audit revealed that the company has configured encryption at rest for ElastiCache However the company did not configure ElastiCache to use encryption in transit Additionally, users can access the cache without authentication A solutions architect must make changes to require user authentication and to ensure that the company is using end-to-end encryption Which solution will meet these requirements?

A. Create an AUTH token Store the token in AWS System Manager Parameter Store, as an encrypted parameter Create a new cluster with AUTH and configure encryption in transit Update the application to retrieve the AUTH token from Parameter Store when necessary and to use the AUTH token for authentication
B. Create an AUTH token Store the token in AWS Secrets Manager Configure the existing cluster to use the AUTH token and configure encryption in transit Update the application to retrieve the AUTH token from Secrets Manager when necessary and to use the AUTH token for authentication.
C. Create an SSL certificate Store the certificate in AWS Secrets Manager Create a new cluster and configure encryption in transit Update the application to retrieve the SSL certificate from Secrets Manager when necessary and to use the certificate for authentication.
D. Create an SSL certificate Store the certificate in AWS Systems Manager Parameter Store, as an encrypted advanced parameter Update the existing cluster to configure encryption in transit Update the application to retrieve the SSL certificate from Parameter Store when necessary and to use the certificate for authentication
Answer: (SHOW ANSWER)
Explanation
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: 149
A company has established a new AWS account. The account is newly provisioned and no changed have been made to the default settings. The company is concerned about the security of the AWS account root user.
What should be done to secure the root user?
A. Create IAM users for daily administrative tasks. Disable the root user.
B. Generate an access key for the root user. Use the access key for daily administration tasks instead of the AWS Management Console.
C. Create IAM users for daily administrative tasks. Enable multi-factor authentication on the root user.

D. Provide the root user credentials to the most senior solutions architect. Have the solutions architect use the root user for daily administration tasks.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 150
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. 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: D (LEAVE A REPLY)
This feature also allows you to connect to any of the participating VPCs from any Direct Connect location, further reducing your costs for making using AWS services on a cross-region basis. https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter-region-vpc-access/ A: There is only a single DC and hence is not highly available.
B: VPC peering means there are additional cost charges when data transfer between region. Also there is a 125 VPC peering limit. Data transferred across Inter-Region VPC Peering connections is charged at the standard inter-region data transfer rates. https://aws.amazon.com/about- aws/whats-new/2017/11/announcing-support-for-inter-region-vpc-peering/ C: Similar to B.
D: Remember one caveat which the question did not state is if there are multiple accounts: The VPCs that reference a particular Direct Connect Gateway must have IP address ranges that do not overlap. Today, the VPCs must all be in the same AWS account; we plan to make this more flexible in the future. https://aws.amazon.com/blogs/aws/new-aws-direct-connect-gateway-inter- region-vpc-access/

NEW QUESTION: 151
A company has an application that generates reports and stores them in an Amazon bucket Amazon S3 bucket.

When a user accesses their report, the application generates a signed URL to allow the user to download the report. The company's security team has discovered that the files are public and that anyone can download them without authentication. The company has suspended the generation of new reports until the problem is resolved. Which set of action will immediately remediate the security issue without impacting the application's normal workflow?

A. Use the Block Public Access feature in Amazon S3 to set the IgnorePublicAcis option to TRUE on the bucket.

B. Review the AWS Trusted advisor bucket permissions check and implement the recommend actions.

C. Run a scrip that puts a Private ACL on all of the object in the bucket.

D. Create an AWS Lambda 'function that applies all policy for users who are not authenticated. Create a scheduled event to invoke the Lambda function.

Answer: (SHOW ANSWER)

NEW QUESTION: 152

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. Which of the following is the short version of the Numeric Condition "NumericLessThanEquals"?

A. numlteq

B. numlteql

C. numltequals

D. numeql

Answer: (SHOW ANSWER)

When using Numeric Conditions within IAM, short versions of the available comparators can be used instead of the more verbose versions. For instance, numlteq is the short version of NumericLessThanEquals.
http://awsdocs.s3.amazonaws.com/SQS/2011-10-01/sqs-dg-2011-10-01.pdf

NEW QUESTION: 153

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 (or 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 markeling 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 SOS queue and trigger an AWS Lambda function lo process them.
C. Receive the orders using the AWS Step Functions program and trigger an Amazon ECS container lo process them.
D. Receive the orders in Amazon Kinesis Data Streams and use Amazon EC2 instances to process them.
Answer: (SHOW ANSWER)
Explanation
Q: How does Amazon Kinesis Data Streams differ from Amazon SQS?
Amazon Kinesis Data Streams enables real-time processing of streaming big data. It provides ordering of records, as well as the ability to read and/or replay records in the same order to multiple Amazon Kinesis Applications. The Amazon Kinesis Client Library (KCL) delivers all records for a given partition key to the same record processor, making it easier to build multiple applications reading from the same Amazon Kinesis data stream (for example, to perform counting, aggregation, and filtering).
https://aws.amazon.com/kinesis/data-streams/faqs/
https://aws.amazon.com/blogs/big-data/unite-real-time-and-batch-analytics-using-the-big-data-lambda-architectu

NEW QUESTION: 154
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 maxmum 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. Set up an Amazon CloudFront distribution for all site contents, and point the distribution at the ALB.
D. Configure an Amazon CloudFront distribution. Point the distribution to an S3 bucket, and migrate the videos from EFS to Amazon S3.
Answer: (SHOW ANSWER)

NEW QUESTION: 155
A company has a large on-premises Apache Hadoop cluster with a 20 PB HDFS database. The cluster is growing every quarter by roughly 200 instances and 1 PB. The company's goals are to enable resiliency for its Hadoop data, limit the impact of losing cluster nodes, and significantly reduce costs. The current cluster runs 24/7 and supports a variety of analysis workloads, including interactive queries and batch processing.
Which solution would meet these requirements with the LEAST expense and down time?
A. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the on-premises cluster. Store the data on EMRFS. Minimize costs using Reserved Instances for master and core nodes and Spot

Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metrics. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
B. Use AWS Snowmobile to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster of similar size and configuration to the current cluster. Store the data on EMRFS.
Minimize costs by using Reserved Instances. As the workload grows each quarter, purchase additional Reserved Instances and add to the cluster.
C. Use AWS Snowball to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workloads based on historical data from the on-premises cluster. Store the on EMRFS. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metrics. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
D. Use AWS Direct Connect to migrate the existing cluster data to Amazon S3. Create a persistent Amazon EMR cluster initially sized to handle the interactive workload based on historical data from the on-premises cluster. Store the data on EMRFS. Minimize costs using Reserved Instances for master and core nodes and Spot Instances for task nodes, and auto scale task nodes based on Amazon CloudWatch metrics. Create job-specific, optimized clusters for batch workloads that are similarly optimized.
Answer: A (LEAVE A REPLY)
Explanation
Q: How should I choose between Snowmobile and Snowball?
To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

NEW QUESTION: 156
You are designing a photo-sharing mobile app. The application will store all pictures in a single Amazon S3 bucket.
Users will upload pictures from their mobile device directly to Amazon S3 and will be able to view and download their own pictures directly from Amazon S3.
You want to configure security to handle potentially millions of users in the most secure manner possible.
What should your server-side application do when a new user registers on the photo-sharing mobile application?
A. Create an IAM user. Update the bucket policy with appropriate permissions for the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
B. Create an IAM user. Assign appropriate permissions to the IAM user. Generate an access key and secret key for the IAM user, store them in the mobile app and use these credentials to access Amazon S3.
C. Create a set of long-term credentials using AWS Security Token Service with appropriate permissions.
Store these credentials in the mobile app and use them to access Amazon S3.
D. Record the user's information in Amazon RDS and create a role in IAM with appropriate permissions.

When the user uses their mobile app, create temporary credentials using the AWS Security Token Service "AssumeRole" function. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
E. Record the user's information in Amazon DynamoDB. When the user uses their mobile app, create temporary credentials using AWS Security Token Service with appropriate permissions. Store these credentials in the mobile app's memory and use them to access Amazon S3. Generate new credentials the next time the user runs the mobile app.
Answer: (SHOW ANSWER)
Explanation
We can use either RDS or DynamoDB, however in our given answers, IAM role is mentioned only with RDS, so I would go with Answer B.
Question was explicitly focused on security, so IAM with RDS is the best choice.

NEW QUESTION: 157
You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2 8xlarge instance type, whose CPUs are mostly idle during processing.
Which of the below would be the most cost efficient way to reduce the runtime of the job?
A. Use smaller instances that have higher aggregate I/O performance.
B. Create more, smaller flies on Amazon S3.
C. Add additional cc2 8xlarge instances by introducing a task group.
D. Create fewer, larger files on Amazon S3.
Answer: (SHOW ANSWER)

NEW QUESTION: 158
A company has 50 AWS accounts that are members of an organization in AWS Organizations Each account contains multiple VPCs The company wants to use AWS Transit Gateway to establish connectivity between the VPCs in each member account Each time a new member account is created, the company wants to automate the process of creating a new VPC and a transit gateway attachment.
Which combination of steps will meet these requirements? (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)

Explanation

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.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-ec2-transitgatewayattachmen

NEW QUESTION: 159

A fitness tracking company serves users around the world, with its primary markets in North America and Asia. The company needs to design an infrastructure for its read-heavy user authorization application with the following requirements:

* Be resilient to problems with the application in any Region.
* Write to a database in a single Region.
* Read from multiple Regions.
* Support resiliency across application tiers in each Region.
* Support the relational database semantics reflected in the application.

Which combination of steps should a solutions architect take? (Select TWO.)

A. Use an Amazon Route 53 geoproximity routing policy combined with a multivalue answer routing policy.

B. Deploy web. application, and MySQL database servers to Amazon EC2 instances in each Region. Set up the application so that reads and writes are local to the Region. Create snapshots of the web, application, and database servers and store the snapshots in an Amazon S3 bucket in both Regions. Set up cross-Region replication for the database layer.

C. Use an Amazon Route 53 geolocation routing policy combined with a failover routing policy.

D. Set up web, application, and Amazon RDS for MySQL instances in each Region. Set up the application so that reads are local and writes are partitioned based on the user. Set up a Multi-AZ failover for the web, application, and database servers. Set up cross-Region replication for the database layer.

E. Set up active-active web and application servers in each Region. Deploy an Amazon Aurora global database with clusters in each Region. Set up the application to use the in-Region Aurora database endpoints. Create snapshots of the web and application servers and store them in an Amazon S3 bucket in both Regions.

Answer: C,E (LEAVE A REPLY)

Explanation

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html Geoproximity routing policy is good to control the user traffic to specific regions. However, a multivalue answer routing policy may cause the users to be randomly sent to other healthy regions that may be far away from the user's location. You can use geolocation routing policy to direct the North American users to your servers on the North America region and configure failover routing to the Asia region in case the North America region fails. You can configure the same for the Asian users pointed to the Asia region servers and have the North America region as its backup.

NEW QUESTION: 160

A company has deployed an application to multiple environments in AWS. including production and testing the company has separate accounts for production and testing, and users are allowed to create additional

application users for team members or services. as needed. The security team has asked the operations team tor better isolation between production and testing with centralized controls on security credentials and improved management of permissions between environments Which of the following options would MOST securely accomplish this goal?

A. Create a script that runs on each account that checks user accounts For adherence to a security policy. Disable any user or service accounts that do not comply.

B. Modify permissions in the production and testing accounts to limit creating new IAM users to members of the operations team Set a strong IAM password policy on each account Create new IAM users and groups in each account to Limit developer access to just the services required to complete their job function.

C. Create a new AWS account to hold user and service accounts, such as an identity account Create users and groups m the identity account. Create roles with appropriate permissions in the production and testing accounts Add the identity account to the trust policies for the roles

D. Create all user accounts in the production account Create roles for access in me production account and testing accounts. Grant cross-account access from the production account to the testing account

Answer: (SHOW ANSWER)

NEW QUESTION: 161

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 PostgreSQL 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 PostgreSQL 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 Amazon RDS 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 PostgreSQL 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.

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.

Answer: D (LEAVE A REPLY)

The correct answer is D.

D). This solution meets the requirements because it uses a scheduled scaling policy for the EC2 instances, which can adjust the capacity according to the known sale events. It also uses Amazon Aurora PostgreSQL Multi-AZ DB cluster, which provides high availability and durability for the database. It uses Amazon EventBridge rules and AWS Lambda functions to create a larger Aurora Replica before a sale event and fail over to it, which can improve the performance and handle the increased traffic. It also uses another EventBridge rule and Lambda function to scale down the Aurora Replica after the sale event, which can save costs123

A). This solution is incorrect because it uses predictive scaling policy for the EC2 instances, which is not suitable for one-time events that are scheduled. Predictive scaling is based on historical data and machine learning, which may not accurately forecast the demand for sale events. It also uses Amazon Aurora PostgreSQL Serverless v2 Multi-AZ DB instance, which does not support read replicas. The use of AWS Step Functions state machine and Lambda functions to pre-warm the database is unnecessary and adds complexity45

B). This solution is incorrect because it uses Amazon RDS for PostgreSQL Multi-AZ DB instance with automatically scaling read replicas, which may not provide enough performance improvement for the sale events. The use of EventBridge rules and Lambda functions to create a larger read replica and fail over to it is risky and may cause downtime or data loss. The use of another EventBridge rule and Lambda function to scale down the read replica is also risky and may cause inconsistency or data loss67

C). This solution is incorrect because it uses predictive scaling policy for the EC2 instances, which is not suitable for one-time events that are scheduled. Predictive scaling is based on historical data and machine learning, which may not accurately forecast the demand for sale events. The use of AWS Step Functions state machine and Lambda functions to pre-warm the database is unnecessary and adds complexity45 References:

1: Scheduled scaling for Amazon EC2 Auto Scaling 2: Amazon Aurora PostgreSQL features 3: Amazon EventBridge rules 4: Predictive scaling for Amazon EC2 Auto Scaling 5: Amazon Aurora Serverless v2 6: Multi-AZ DB instance deployments - Amazon Relational Database Service 7: Working with PostgreSQL read replicas - Amazon Relational Database Service

NEW QUESTION: 162

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: 163
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 Conng rule that requires all resources in the selected OUs to associate the AWS WAF rules. Deploy automated remediation actions by using AWS Lambda to fix noncompliant resources. Deploy AWS WAF rules by using an AWS CloudFormation stack set to target the same OUs where the AWS Config rule is applied.
C. Create AWS WAF rules in the management account of the organization. Use AWS Lambda environment variables to store account numbers and OUs to manage Update environment variables as needed to add or remove accounts or OUs Create cross-account IAM roles in member accounts. Assume the roles by using AWS Security Token Service (AWS STS) in the Lambda function to create and update AWS WAF rules in the member accounts
D. Use AWS Control Tower to manage AWS WAF rules across accounts in the organization. Use AWS Key Management Service (AWS KMS) to store account numbers and OUs to manage Update AWS KMS as needed to add or remove accounts or OUs. Create IAM users in member accounts Allow AWS Control Tower in the management account to use the access key and secret access key to create and update AWS WAF rules in the member accounts

Answer: A (LEAVE A REPLY)

Explanation

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: 164

Your team has a tomcat-based Java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis.

Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC .The optimal setup for persistence and security that meets the above requirements would be the following.

A. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable Alter its security group to allow access to It from hosts In your application subnets.

B. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable. Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.

C. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.

D. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code Alter its security group to allow access to it from hosts within your VPC's IP address block.

Answer: (SHOW ANSWER)

NEW QUESTION: 165

With Amazon Elastic MapReduce (Amazon EMR) you can analyze and process vast amounts of data. The cluster is managed using an open-source framework called Hadoop.

You have set up an application to run Hadoop jobs. The application reads data from DynamoDB and generates a temporary file of 100 TBs.

The whole process runs for 30 minutes and the output of the job is stored to S3. Which of the below mentioned options is the most cost effective solution in this case?

A. Use Spot Instances to run Hadoop jobs and configure them with EBS volumes for persistent data storage.

B. Use Spot Instances to run Hadoop jobs and configure them with ephermal storage for output file storage.

C. Use an on demand instance to run Hadoop jobs and configure them with EBS volumes for persistent

storage.
D. Use an on demand instance to run Hadoop jobs and configure them with ephemeral storage for output file storage.
Answer: B (LEAVE A REPLY)
AWS EC2 Spot Instances allow the user to quote his own price for the EC2 computing capacity. The user can simply bid on the spare Amazon EC2 instances and run them whenever his bid exceeds the current Spot Price. The Spot Instance pricing model complements the On-Demand and Reserved Instance pricing models, providing potentially the most cost-effective option for obtaining compute capacity, depending on the application. The only challenge with a Spot Instance is data persistence as the instance can be terminated whenever the spot price exceeds the bid price.
In the current scenario a Hadoop job is a temporary job and does not run for a longer period. It fetches data from a persistent DynamoDB. Thus, even if the instance gets terminated there will be no data loss and the job can be re-run. As the output files are large temporary files, it will be useful to store data on ephermal storage for cost savings.
Reference: http://aws.amazon.com/ec2/purchasing-options/spot-instances/

NEW QUESTION: 166
A company has built a high performance computing (HPC) cluster in AWS for a tightly coupled workload that generates a large number of shared files stored in Amazon EFS. The cluster was performing well when the number of Amazon EC2 instances in the cluster was 100. However, when the company increased the cluster size to 1,000 EC2 instances, overall performance was well below expectations Which collection of design choices should a solutions architect make to achieve the maximum performance from the HPC cluster? (Select THREE.)
A. Replace Amazon EFS with Amazon FSx for Lustre.
B. Ensure the cluster is launched across multiple Availability Zones.
C. Select EC2 instance types with an Elastic Fabric Adapter (EFA) enabled
D. Launch the EC2 instances and attach elastic network interfaces in multiples of four.
E. Replace Amazon EFS with multiple Amazon EBS volumes in a RAID array.
F. Ensure the HPC cluster is launched within a single Availability Zone.
Answer: C,E,F (LEAVE A REPLY)

NEW QUESTION: 167
Your team has a tomcat-based java application you need to deploy into development, test and production environments. After some research, you opt to use Elastic Beanstalk due to its tight integration with your developer tools and RDS due to its ease of management. Your QA team lead points out that you need to roll a sanitized set of production data into your environment on a nightly basis. Similarly, other software teams in your org want access to that same restored data via their EC2 instances in your VPC. The optimal setup for persistence and security that meets the above requirements would be the following:
A. Create your RDS instance separately and add its IP address to your application's DB connection strings in your code.
Alter its security group to allow access to it from hosts within your VPC's IP address block.
B. Create your RDS instance separately and pass its DNS name to your's DB connection string as an environment variable.
Alter its security group to allow access to it from hosts in your application subnets.
C. Create your RDS instance as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
D. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variable.
Create a security group for client machines and add it as a valid source for DB traffic to the security group of the RDS instance itself.
Answer: C (LEAVE A REPLY)
Elastic Beanstalk provides support for running Amazon RDS instances in your Elastic Beanstalk environment. This works great for development and testing environments, but is not ideal for a production environment because it ties the lifecycle of the database instance to the lifecycle of your application's environment.
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo.RDS.html

NEW QUESTION: 168
A photo-sharing and publishing company receives 10,000 to 150,000 images daily. The company receives the images from multiple suppliers and users registered with the service. The company is moving to AWS and wants to enrich the existing metadata by adding data using Amazon Rekognition.
The following is an example of the additional data:
```
list celebrities [name of the personality] wearing [color] looking [happy, sad] near [location example Eiffel Tower in Paris]
```
As part of the cloud migration program, the company uploaded existing image data to Amazon S3 and told users to upload images directly to Amazon S3.
What should the Solutions Architect do to support these requirements?
A. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognition. Use Amazon DynamoDB to store the metadata and Amazon ES to create an index. Use a web front-end to provide search capabilities backed by Amazon ES.
B. Use Amazon Kinesis to stream data based on an S3 event. Use an application running in Amazon EC2 to extract metadata from the images. Then store the data on Amazon DynamoDB and Amazon CloudSearch and create an index. Use a web front-end with search capabilities backed by CloudSearch.

C. Start an Amazon SQS queue based on S3 event notifications. Then have Amazon SQS send the metadata information to Amazon DynamoDB. An application running on Amazon EC2 extracts data from Amazon Rekognition using the API and adds data to DynamoDB and Amazon ES. Use a web front-end to provide search capabilities backed by Amazon ES.
D. Trigger AWS Lambda based on an S3 event notification to create additional metadata using Amazon Rekognition. Use Amazon RDS MySQL Multi-AZ to store the metadata information and use Lambda to create an index. Use a web front-end with search capabilities backed by Lambda.
Answer: A (LEAVE A REPLY)
https://github.com/aws-samples/lambda-refarch-imagerecognition

NEW QUESTION: 169
A company runs a public-facing application that uses a Java-based web sen/ice via a RESTful API It is hosted on Apache Tomcat on a single server in a data center that runs consistently at 30% CPU utilization Use of the API is expected to increase by 10 times with a new product launch The business wants to migrate the application to AWS with no disruption and needs it to scale to meet demand The company has already decided to use Amazon Route 53 and CNAME records lo redirect traffic How can these requirements be met with the LEAST amount of effort?
A. Create a Docker image and migrate the image to Amazon ECS Then change the application code to direct web service queries to the ECS container
B. Use AWS Elastic Beanstalk to deploy the Java web service and enable Auto Scaling Then switch the application to use the new web service
C. Modify the application to call the web service via Amazon API Gateway Then create a new AWS Lambda Java function to run the Java web service code After testing change API Gateway to use the Lambda function
D. Lift and shift the Apache server to the cloud using AWS SMS Then switch the application to direct web service traffic to the new instance
Answer: B (LEAVE A REPLY)

NEW QUESTION: 170
In the context of AWS IAM, identify a true statement about user passwords (login profiles).
A. They must contain Unicode characters.
B. They can contain any Basic Latin (ASCII) characters.
C. They must begin and end with a forward slash (/).
D. They cannot contain Basic Latin (ASCII) characters.
Answer: B (LEAVE A REPLY)
The user passwords (login profiles) of IAM users can contain any Basic Latin (ASCII) characters.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION: 171
By default, Amazon Cognito maintains the last-written version of the data.
You can override this behavior and resolve data conflicts programmatically.

In addition, push synchronization allows you to use Amazon Cognito to send a silent _____ notification to all devices associated with an identity to notify them that new data is available.

A. get

B. post

C. pull

D. push

Answer: (SHOW ANSWER)

http://aws.amazon.com/cognito/faqs/

NEW QUESTION: 172

_____ pricing offers significant savings over the normal price of DynamoDB provisioned throughput capacity.

A. Discount Voucher

B. Reserved Capacity

C. Discount Service

D. Reserved Point

Answer: B (LEAVE A REPLY)

Reserved Capacity pricing offers significant savings over the normal price of DynamoDB provisioned throughput capacity. When you buy Reserved Capacity, you pay a one-time upfront fee and commit to paying for a minimum usage level, at the hourly rates indicated above, for the duration of the Reserved Capacity term.

http://aws.amazon.com/dynamodb/pricing/

NEW QUESTION: 173

A development team is collaborating with another company to create an integrated product. The other company needs to access an Amazon Simple Queue Service (Amazon SQS) queue that is contained in the development team's account. The other company wants to poll the queue without giving up its own account permissions to do so.

How should a solutions architect provide access to the SQS queue?

A. Create an IAM policy that provides the other company access to the SQS queue.

B. Create an SQS access policy that provides the other company access to the SQS queue.

C. Create an Amazon Simple Notification Service (Amazon SNS) access policy that provides the other company access to the SQS queue.

D. Create an instance profile that provides the other company access to the SQS queue.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 174

A user has configured an EC2 instance in the US-East-1a zone. The user has enabled detailed monitoring of the instance. The user is trying to get the data from CloudWatch using a CLI. Which of the below mentioned CloudWatch endpoint URLs should the user use?

A. monitoring.us-east-1a.amazonaws.com

B. cloudwatch.us-east-1a.amazonaws.com

C. monitoring.us-east-1.amazonaws.com
D. monitoring.us-east-1-a.amazonaws.com
Answer: C (LEAVE A REPLY)
The CloudWatch resources are always region specific and they will have the end point as region specific. If the user is trying to access the metric in the US-East-1 region, the endpoint URL will be: monitoring.us-east-1.amazonaws.com
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/regions_endpoints.html

NEW QUESTION: 175
A solutions architect is performing a security review of a recently migrated workload. The workload is a web application that consists of Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer. The solutions architect must improve the security posture and minimize the impact of a DDoS attack on resources.
Which solution is MOST effective?
A. Enable Amazon GuardDuty and configure findings written to Amazon CloudWatch. Create an event with CloudWatch Events for DDoS alerts that triggers Amazon Simple Notification Service (Amazon SNS).
Have Amazon SNS invoke a custom AWS Lambda function that parses the logs, looking for a DDoS attack. Modify a network ACL to block identified source IP addresses.
B. Enable VPC Flow Logs and store then in Amazon S3. Create a custom AWS Lambda functions that parses the logs looking for a DDoS attack. Modify a network ACL to block identified source IP addresses.
C. Configure an AWS WAF ACL with rate-based rules. Create an Amazon CloudFront distribution that points to the Application Load Balancer. Enable the WAF ACL on the CloudFront distribution.
D. Create a custom AWS Lambda function that adds identified attacks into a common vulnerability pool to capture a potential DDoS attack. Use the identified information to modify a network ACL to block access.
Answer: (SHOW ANSWER)

NEW QUESTION: 176
A solutions architect is auditing the security setup of an AWS Lambda function for a company. The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function.
The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.
What should the solutions architect recommend to meet these requirements?
A. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

B. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Enforce HTTPS on the connection to Amazon S3 during data transfers.
C. Save the database credentials in AWS Systems Manager Parameter Store. Set up password rotation on the credentials in Parameter Store. Change the IAM role for the Lambda function to allow the function to access Parameter Store. Modify the Lambda function to retrieve the credentials from Parameter Store.
Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
D. Save the database credentials in AWS Secrets Manager. Set up password rotation on the credentials in Secrets Manager. Change the IAM role for the Lambda function to allow the function to access Secrets Manager. Modify the Lambda function to retrieve the credentials Om Secrets Manager. Enforce HTTPS on the connection to Amazon S3 during data transfers.
Answer: (SHOW ANSWER)
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.html

NEW QUESTION: 177
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: (SHOW ANSWER)
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: 178
A company is moving a business-critical, multi-tier application to AWS. The architecture consists of a desktop client application and server infrastructure. The server infrastructure resides in an on-premises data center that

frequently fails to maintain the application uptime SLA of 99.95%. A Solutions Architect must re-architect the application to ensure that it can meet or exceed the SLA.
The application contains a PostgreSQL database running on a single virtual machine. The business logic and presentation layers are load balanced between multiple virtual machines. Remote users complain about slow load times while using this latency-sensitive application.
Which of the following will meet the availability requirements with little change to the application while improving user experience and minimizing costs?
A. Migrate the database to an Amazon Redshift cluster with at least two nodes. Combine and host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balancer. Use Amazon CloudFront to improve the user experience.
B. Migrate the database to a PostgreSQL database in Amazon EC2. Host the application and presentation layers in automatically scaled Amazon ECS containers behind an Application Load Balancer. Allocate an Amazon WorkSpaces WorkSpace for each end user to improve the user experience.
C. Migrate the database to an Amazon RDS PostgreSQL Multi-AZ configuration. Host the application and presentation layers in automatically scaled AWS Fargate containers behind a Network Load Balancer. Use Amazon ElastiCache to improve the user experience.
D. Migrate the database to an Amazon RDS Aurora PostgreSQL configuration. Host the application and presentation layers in an Auto Scaling configuration on Amazon EC2 instances behind an Application Load Balancer. Use Amazon AppStream 2.0 to improve the user experience.
Answer: (SHOW ANSWER)

NEW QUESTION: 179
A company has an image processing workload running on Amazon Elastic Container Service (Amazon ECS) in two private subnets. Each private subnet uses a NAT instance for internet access. All images are stored in Amazon S3 buckets. The company is concerned about the data transfer costs between Amazon ECS and Amazon S3.
What should a solutions architect do to reduce costs?
A. Configure a gateway endpoint for traffic destined to Amazon S3.
B. Configure Amazon CloudFront for the S3 bucket storing the images.
C. Configure a NAT gateway to replace the NAT instances.
D. Configure an interface endpoint for traffic destined to Amazon S3.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 180
A solutions architect plans to convert a company's monolithic web application into a multi-tier application.
The company wants to avoid managing its own infrastructure. The minimum requirements for the web application are high availability, scalability, and regional low latency during peak hours. The solution should also store and retrieve data with millisecond latency using the application's API.
Which solution meets these requirements?
A. Use AWS Fargate to host the web application with backend Amazon RDS Multi-AZ DB instances.

B. Use Amazon API Gateway with an edge-optimized API endpoint, AWS Lambda for compute, and Amazon DynamoDB as the data store.
C. Use an Amazon Route 53 routing policy with geolocation that points to an Amazon S3 bucket with static website hosting and Amazon DynamoDB as the data store.
D. Use an Amazon CloudFront distribution that points to an Elastic Load Balancer with an Amazon EC2 Auto Scaling group, along with Amazon RDS Multi-AZ DB instances.
Answer: (SHOW ANSWER)
Explanation/Reference:

NEW QUESTION: 181
A company is implementing a multi-account strategy; however, the Management team has expressed concerns that services like DNS may become overly complex. The company needs a solution that allows private DNS to be shared among virtual private clouds (VPCs) in different accounts. The company will have approximately 50 accounts in total.
What solution would create the LEAST complex DNS architecture and ensure that each VPC can resolve all AWS resources?
A. Create a shared services VPC in a central account, and create a VPC peering connection from the shared services VPC to each of the VPCs in the other accounts. Within Amazon Route 53, create a privately hosted zone in the shared services VPC and resource record sets for the domain and subdomains. Programmatically associate other VPCs with the hosted zone.
B. Create a VPC peering connection among the VPCs in all accounts. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "true" for each VPC. Create an Amazon Route 53 private zone for each VPC. Create resource record sets for the domain and subdomains. Programmatically associate the hosted zones in each VPC with the other VPCs.
C. Create a shared services VPC in a central account. Create a VPC peering connection from the VPCs in other accounts to the shared services VPC. Create an Amazon Route 53 privately hosted zone in the shared services VPC with resource record sets for the domain and subdomains. Allow UDP and TCP port 53 over the VPC peering connections.
D. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "false" in every VPC. Create an AWS Direct Connect connection with a private virtual interface. Allow UDP and TCP port 53 over the virtual interface. Use the on-premises DNS servers to resolve the IP addresses in each VPC on AWS.
Answer: (SHOW ANSWER)
https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-with-amazon-route-53-and-aws-transit-gateway/

NEW QUESTION: 182
A company will several AWS accounts is using AWS Organizations and service control policies (SCPs). An Administrator created the following SCP and has attached it to an organizational unit (OU) that contains AWS account 1111-1111-1111:

Developers working in account 1111-1111-1111 complain that they cannot create Amazon S3 buckets. How should the Administrator address this problem?
A. Add s3:CreateBucket with "Allow" effect to the SCP.
B. Remove the SCP from account 1111-1111-1111.
C. Instruct the Developers to add Amazon S3 permissions to their IAM entities.
D. Remove the account from the OU, and attach the SCP directly to account 1111-1111-1111.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 183
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 partition placement group
B. Launch compute optimized EC2 instances in a partition placement group
C. Launch memory optimized EC2 instances in a cluster placement group
D. Launch compute optimized EC2 instances in a spread placement group.
Answer: C (LEAVE A REPLY)
Explanation
According to the AWS Certified Solutions Architect - Professional Official Amazon Text Book, a cluster placement group is a logical grouping of instances within a single Availability Zone. It is recommended for applications that need low network latency, high network throughput, or both. Network latency and throughput are improved because instances within the same placement group can communicate directly with each other over a high-speed, low-latency network connection. Memory-optimized instances are ideal for in-memory databases as they are optimized for high performance and low latency.

NEW QUESTION: 184
A Solutions Architect must establish a patching plan for a large mixed fleet of Windows and Linux servers. The patching plan must be implemented securely, be audit ready, and comply with the company's business requirements.
Which option will meet these requirements with MINIMAL effort?
A. Install and use an OS-native patching service to manage the update frequency and release approval for all instances. Use AWS Config to verify the OS state on each instance and report on any patch compliance issues.
B. Use AWS Systems Manager on all instances to manage patching. Test patches outside of production and then deploy during a maintenance window with the appropriate approval.
C. Use AWS OpsWorks for Chef Automate to run a set of scripts that will iterate through all instances of a given type. Issue the appropriate OS command to get and install updates on each instance, including any required restarts during the maintenance window.
D. Migrate all applications to AWS OpsWorks and use OpsWorks automatic patching support to keep the OS up-to-date following the initial installation. Use AWS Config to provide audit and compliance reporting.
Answer: B (LEAVE A REPLY)
Only Systems Manager can patch both OS effectively on AWS and on premise.

NEW QUESTION: 185
A sys admin is maintaining an application on AWS. The application is installed on EC2 and user has configured ELB and Auto Scaling. Considering future load increase, the user is planning to launch new servers proactively so that they get registered with ELB.
How can the user add these instances with Auto Scaling?
A. Decrease the minimum limit of the Auto Scaling group
B. Increase the maximum limit of the Auto Scaling group
C. Launch an instance manually and register it with ELB on the fly
D. Increase the desired capacity of the Auto Scaling group
Answer: D (LEAVE A REPLY)
Explanation

A user can increase the desired capacity of the Auto Scaling group and Auto Scaling will launch a new instance as per the new capacity. The newly launched instances will be registered with ELB if Auto Scaling group is configured with ELB. If the user decreases the minimum size the instances will be removed from Auto Scaling. Increasing the maximum size will not add instances but only set the maximum instance cap.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-manual-scaling.html

NEW QUESTION: 186
A company is building an image service on the web that will allow users to upload and search random photos. At peak usage, up to 10.000 users worldwide will upload their images. The service will then overlay text on the uploaded images, which will then be published on the company website.
Which design should a solutions architect implement?
A. Store the uploaded images on a shared Amazon Elastic Block Store (Amazon EBS) volume mounted to a fleet of Amazon EC2 Spot instances. Create an Amazon DynamoDB table that contains information about each uploaded image and whether it has been processed. Use an Amazon EventBridge (Amazon CloudWatch Events) rule lo scale out EC2 instances. Enable Amazon CloudFront and configure the origin to reference an Elastic Load Balancer in front of the fleet of EC2 instances.
B. Store the uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to the Amazon Simple Queue Service (Amazon SOS) queue. Create a fleet of Amazon EC2 instances to pull messages from Ihe SOS queue to process the images and place them in another S3 bucket. Use Amazon CloudWatch metrics for queue depth to scale out EC2 instances. Enable Amazon CloudFront and configure the origin to be the S3 bucket that contains the processed images.
C. Store Ihe uploaded images in an Amazon S3 bucket and configure an S3 bucket event notification to send a message to Amazon Simple Notification Service (Amazon SNS). Create a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) to pull messages from Amazon SNS to process the images and place them in Amazon Elastic File System (Amazon EFS). Use Amazon CloudWatch metrics for the SNS message volume to scale out EC2 instances. Enable Amazon CloudFront and configure the origin lo be the ALB in front of the EC2 instances.
D. Store Ihe uploaded images in Amazon Elastic File System (Amazon EFS). Send application log information about each image to Amazon CloudWatch Logs. Create a fleet of Amazon EC2 instances that use CloudWatch Logs to determine which images need to be processed. Place processed images in anolher directory in Amazon EFS. Enable Amazon CloudFront and configure the origin to be Ihe one of the EC2 instances in the fleet.
Answer: (SHOW ANSWER)

NEW QUESTION: 187
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 Region solution that wifi reduce latency improve reliability, and require the least effort to implement What should the solutions architect do to meet these requirements'

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

NEW QUESTION: 188
The user has provisioned the PIOPS volume with an EBS optimized instance. Generally speaking, in which I/O chunk should the bandwidth experienced by the user be measured by AWS?
A. 128 KB
B. 256 KB
C. 64 KB
D. 32 KB
Answer: B (LEAVE A REPLY)
IOPS are input/output operations per second. Amazon EBS measures each I/O operation per second (that is 256 KB or smaller) as one IOPS.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

NEW QUESTION: 189
A company stores customer data in an Amazon S3 bucket with S3 Versioning enabled in the us>west-2 Region. The S3 bucket is encrypted with an AWS Key Management Service (AWS KMS) customer managed CMK A compliance policy states that redundant copies of all S3 objects must be stored in the us-east-2 Region The S3 buckets are allowed to stay in the same AWS account Which combination of steps will meet these requirements with the LEAST operational effort? (Select THREE)
A. Create a destination S3 bucket in us-east-2 with S3 Versioning enabled
B. Create and assign an S3 bucket policy that allows reading from the source S3 bucket.
C. Create a destination S3 bucket in us-east-2
D. Configure an AWS Lambda function that copies objects to the us-east-2 bucket and is triggered when objects are created in the us-west-2 bucket
E. Create and assign to Amazon S3 an IAM role with a policy that allows reading from the source S3 bucket and replication to the destination S3 bucket

F. Set up S3 Cross-Region Replication between the two S3 buckets.
Answer: A,B,F (LEAVE A REPLY)

NEW QUESTION: 190
A company wants to allow its Marketing team to perform SQL queries on customer records to identity market segments. The data is spread across hundred of flies. The records must be encrypted in transit and at rest. The Team Manager must have the ability to manage users and groups users and groups, but no team members should have access or resources not required for the SQL queries. Additionally Administrators needed to audit the queries made and receive notifications when a query violets rules defined by the Security team.
AWS organization has been used to create a new account and an AWS 1AM user with administrator permissions to the team manager.
Which design meets these requirements?
A. Apply a service control policy (SCP) that allows access to 1AM. Amazon RDS. and AWS Cloudtrail Load customer records in Amazon MySQL and train users to execute queries using the AWS CLI Stream the query logs lo Amazon CloudWatch Logs from the RDS database instance. Use a subscription filter with AWS Lambda functions to audit and alarm on queries against personal data.
B. Apply a service control policy (SCP) that denies access to all services except IAM Amazon DynamoDB and AWS CloudTrail . Store custom records in DynamoDB and tram users to execute queries using the AWS CLI Enable DynamoDB queries that are issued and use an AWS Lambda function for real-time monitoring and alerting.
C. Apply a service control policy (SCP) that allow access to AIM, Amazon S3, and AWS CloudTrail. Store customer records as files in Amazon S3 and trail users to leverage the Amazon S3 Select feature and execute queries using the AWS CLI Enable S3 object-level and analyze CloudTrail events to audit and alarm on queries against personal data.
D. Apply a service control policy (SCP) that denies access to all services except IAM Amazon Athena Amazon S3 and AWS CloudTrail Store customer record files in Amazon S3 and train users to execute queries using the CLI via Athena Analyze CloudTrail events lo audit and alarm on queries against personal data
Answer: D (LEAVE A REPLY)

NEW QUESTION: 191
A company wants to run a custom network analysis software package to inspect traffic as traffic leaves and enters a VPC. The company has deployed the solution by using AWS Cloud Formation on three Amazon EC2 instances in an Auto Scaling group. All network routing has been established to direct traffic to the EC2 instances.
Whenever the analysis software stops working, the Auto Scaling group replaces an instance. The network routes are not updated when the instance replacement occurs.
Which combination of steps will resolve this issue? {Select THREE.)
A. Update the Cloud Formation template to install the Amazon CloudWatch agent on the EC2 instances. Configure the CloudWatch agent to send process metrics for the application.
B. In the Cloud Formation template, write a condition that updates the network routes when a replacement instance is launched.

C. Create alarms based on EC2 status check metrics that will cause the Auto Scaling group to replace the failed instance.
D. Create an alarm for the custom metric in Amazon CloudWatch for the failure scenarios. Configure the alarm to publish a message to an Amazon Simple Notification Service {Amazon SNS) topic.
E. Create an AWS Lambda function that responds to the Amazon Simple Notification Service (Amazon SNS) message to take the instance out of service. Update the network routes to point to the replacement instance.
F. Update the Cloud Formation template to install AWS Systems Manager Agent on the EC2 instances. Configure Systems Manager Agent to send process metrics for the application.
Answer: A,D,E (LEAVE A REPLY)

NEW QUESTION: 192
A software as a service (SaaS) company offers a cloud solution for document management to private law firms and the public sector. A Local Government client recently mandated that highly confidential documents cannot be stored outside the country. The company CIO asks a solutions architect to ensure the application can adapt to this new requirement The CIO also wants to have a proper backup plan for these documents, as backups are not currently performed What solution meets these requirements?
A. Tag documents as either regular or secret in Amazon S3. Create an individual S3 backup bucket in the same AWS account and AWS Region. Use S3 selective cross-region replication based on object tags to move regular documents to an S3 bucket in a different AWS Region. Configure an AWS Lambda function that triggers when new S3 objects are created in the main bucket to replicate only documents tagged as secret into the S3 bucket in the same AWS Region
B. Tag documents that are not highly confidential as regular in Amazon S3. Create individual S3 buckets for each user Upload objects to each user's bucket. Set S3 bucket replication from these buckets to a central S3 bucket in a different AWS account and AWS Region. Configure an AWS Lambda function triggered by scheduled events in Amazon CloudWatch to delete objects that are tagged as secret in the S3 backup bucket.
C. Tag highly confidential documents as secret in Amazon S3. Create an individual S3 backup bucket m the same AWS account and AWS Region. Use S3 selective cross-region replication based on object tags to move regular documents to a different AWS Region Create an Amazon CloudWatch Events rule for a S3 objects tagged as secret to trigger an AWS Lambda function to replicate them into a separate bucket in the same AWS Region.
D. Tag documents as either regular or secret in Amazon S3. Create an individual S3 backup bucket in the same AWS account and AWS Region. Create a cross-region S3 bucket in a separate AWS account. Set proper 1AM roles to allow cross-region permissions to the S3 buckets. Configure an AWS Lambda function triggered by Amazon CloudWatch scheduled events to copy objects that are tagged as secret to the S3 backup bucket and objects tagged as normal to the cross-region S3 bucket
Answer: (SHOW ANSWER)

NEW QUESTION: 193
A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC.
How should they architect their solution to achieve these goals?

A. Configure servers running in the VPC using the host-based 'route' commands to send all traffic through the platform to a scalable virtualized IDS/IPS.
B. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform for inspection.
C. Create a second VPC and route all traffic from the primary application VPC through the second VPC where the scalable virtualized IDS/IPS platform resides.
D. Configure an instance with monitoring software and the elastic network interface (ENI) set to promiscuous mode packet sniffing to see a traffic across the VPC.
Answer: (SHOW ANSWER)

NEW QUESTION: 194
A company has more than 100 AWS accounts, with one VPC per account, that need outbound HTTPS connectivity to the internet. The current design contains one NAT gateway per Availability Zone (AZ) in each VPC. To reduce costs and obtain information about outbound traffic, management has asked for a new architecture for internet access.
Which solution will meet the current needs, and continue to grow as new accounts are provisioned, while reducing costs?
A. Create a proxy fleet in a central VPC account. Create an AWS PrivateLink endpoint service in the central VPC. Use PrivateLink interface for internet connectivity through the proxy fleet.
B. Create a central VPC for outbound internet traffic. Use VPC peering to default route to a set of redundant NAT gateway in the central VPC.
C. Create a transit VPC across two AZs using a third-party routing appliance. Create a VPN connection to each VPC. Default route internet traffic to the transit VPC.
D. Create multiple hosted-private AWS Direct Connect VIFs, one per account, each with a Direct Connect gateway. Default route internet traffic back to an on-premises router to route to the internet.
Answer: (SHOW ANSWER)

NEW QUESTION: 195
A company runs an e-commerce platform with front-end and e-commerce tiers. Both tiers run on LAMP stacks with the front-end instances running behind a load balancing appliance that has a virtual offering on AWS. Currently, the operations team uses SSH to log in to the instances to maintain patches and address other concerns. The platform has recently been the target of multiple attacks, including
* A DDoS attack
* An SOL injection attack
* Several successful dictionary attacks on SSH accounts on the web servers.
The company wants to improve the security of the e-commerce platform by migrating to AWS. The company's solutions architects have decided to use the following approach:
* Code review the existing application and fix any SQL injection issues
* Migrate the web application to AWS and leverage the latest AWS Linux AMI to address initial security patching
* Install AWS Systems Manager to manage patching and allow the system administrators to run commands on all instances, as needed.

What additional steps will address at of the identified attack types while providing high availability and minimizing risk?

A. Enable SSH access to the Amazon EC2 instances through a bastion host secured by limiting access to specific IP addresses Migrate on-premises MySQL to a self-managed EC2 instance Leverage an AWS Elastic Load Balancer to spread the load, and enable AWS Shield Standard for DDoS protection. Add an Amazon CloudFront distribution in front of the website.

B. Enable SSH access to the Amazon EC2 instances using a security group that limits access to specific IPs Migrate on-premises MySQL to Amazon RDS Multi-AZ. Install the third-party load balancer from the AWS Marketplace and migrate the existing rules to the load balancer's AWS instances. Enable AWS Shield Standard for DDoS protection.

C. Disable SSH access to the Amazon EC2 instances. Migrate on-premises MySQL to Amazon RDS Multi-AZ Leverage an Elastic Load Balancer to spread the load and enable AWS Shield Advanced for protection. Add an Amazon CloudFront distnbution in front of the website. Enable AWS WAF on the distribution to manage the rules.

D. Disable SSH access to the EC2 instances. Migrate on-premises MySQL to Amazon RDS Single-AZ. Leverage an AWS Elastic Load Balancer to spread the load Add an Amazon CloudFront distribution in front of the website Enable AWS WAF on the distribution to manage the rules

Answer: C (LEAVE A REPLY)

NEW QUESTION: 196

A company is migrating some of its applications to AWS. The company wants to migrate and modernize the applications quickly after it finalizes networking and security strategies. The company has set up an AWS Direct Connection connection in a central network account.

The company expects to have hundreds of AWS accounts and VPCs in the near future. The corporate network must be able to access the resources on AWS seamlessly and also must be able to communicate with all the VPCs. The company also wants to route its cloud resources to the internet through its on-premises data center.

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

A. Provision only private subnets. Open the necessary route on the transit gateway and customer gateway to allow outbound internet traffic from AWS to flow through NAT services that run in the data center.

B. Provision an internet gateway. Attach the internet gateway to subnets. Allow internet traffic through the gateway.

C. Share the transit gateway with other accounts. Attach VPCs to the transit gateway.

D. Create a Direct Connect gateway and a transit gateway in the central network account. Attach the transit gateway to the Direct Connect gateway by using a transit VIF.

E. Create a Direct Connect gateway in the central account. In each of the accounts, create an association proposal by using the Direct Connect gateway and the account ID for every virtual private gateway.

F. Provision VPC peering as necessary.

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

NEW QUESTION: 197
A Development team is deploying new APIs as serverless applications within a company. The team is currently using the AWS Management Console to provision Amazon API Gateway, AWS Lambda, and Amazon DynamoDB resources. A Solutions Architect has been tasked with automating the future deployments of these serverless APIs.
How can this be accomplished?
A. Use AWS CloudFormation to define the serverless application. Implement versioning on the Lambda functions and create aliases to point to the versions. When deploying, configure weights to implement shifting traffic to the newest version, and gradually update the weights as traffic moves over.
B. Use the AWS Serverless Application Model to define the resources. Upload a YAML template and application files to the code repository. Use AWS CodePipeline to connect to the code repository and to create an action to build using AWS CodeBuild. Use the AWS CloudFormation deployment provider in CodePipeline to deploy the solution.
C. Use AWS CloudFormation with a Lambda-backed custom resource to provision API Gateway. Use the AWS::DynamoDB::Table and AWS::Lambda::Function resources to create the Amazon DynamoDB table and Lambda functions. Write a script to automate the deployment of the CloudFormation template.
D. Commit the application code to the AWS CodeCommit code repository. Use AWS CodePipeline and connect to the CodeCommit code repository. Use AWS CodeBuild to build and deploy the Lambda functions using AWS CodeDeploy. Specify the deployment preference type in CodeDeploy to gradually shift traffic over to the new version.
Answer: (SHOW ANSWER)

NEW QUESTION: 198
A company is deploying a third-party firewall appliance solution from AWS Marketplace to monitor and protect traffic that leaves the company's AWS environments. The company wants to deploy this appliance into a shared services VPC and route all outbound internet-bound traffic through the appliances.
A solutions architect needs to recommend a deployment method that prioritizes reliability and minimizes failover time between firewall appliances within a single AWS Region. The company has set up routing from the shared services VPC to other VPCs.
Which steps should the solutions architect recommend to meet these requirements? (Select THREE)
A. Deploy two firewall appliances into the shared services VPC. each in a separate Availability Zone

B. Create a new Gateway Load Balancer in the shared services VPC Create a new target group, and attach it to the new Gateway Load Balancer Add each of the firewall appliance instances to the target group
C. Deploy two firewall appliances into the shared services VPC. each in the same Availability Zone
D. Create a VPC interface endpoint Add a route to the route table in the shared services VPC. Designate the new endpoint as the next hop for traffic that enters the shared services VPC from other VPCs.
E. Create a new Network Load Balancer in the shared services VPC Create a new target group, and attach it to the new Network Load Balancer Add each of the firewall appliance instances to the target group.
Answer: (SHOW ANSWER)

NEW QUESTION: 199
A company is reviewing a recent migration of a three-tier application to a VPC. The security team discovers that the principle of least privilege is not being applied to Amazon EC2 security group ingress and egress rules between the application tiers.
What should a solutions architect do to correct this issue?
A. Create security group rules using the security group ID as the source or destination.
B. Create security group rules using the subnet CIDR blocks as the source or destination.
C. Create security group rules using the VPC CIDR blocks as the source or destination.
D. Create security group rules using the instance ID as the source or destination.
Answer: (SHOW ANSWER)

NEW QUESTION: 200
A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.
Which solution will meet these requirements while providing the FASTEST storage performance?
A. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster. Configure the ReplicaSet to mount the file system. Direct the application to store files in the file system. Configure AWS Backup to back up and retain copies of the data for 1 year.
B. Create an Amazon Elastic Block Store (Amazon EBS) volume. Enable the EBS Multi-Attach feature. Configure the ReplicaSet to mount the EBS volume. Direct the application to store files in the EBS volume. Configure AWS Backup to back up and retain copies of the data for 1 year.
C. Create an Amazon S3 bucket. Configure the ReplicaSet to mount the S3 bucket. Direct the application to store files in the S3 bucket. Configure S3 Versioning to retain copies of the data. Configure an S3 Lifecycle policy to delete objects after 1 year.
D. Configure the ReplicaSet to use the storage available on each of the running application pods to store the files locally. Use a third-party tool to back up the EKS cluster for 1 year.
Answer: A (LEAVE A REPLY)
Explanation
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: 201
A bicycle sharing company is developing a multi-tier architecture to track the location of its bicycles during peak operating hours. The company wants to use these data points in its existing analytics platform. A solutions architect must determine the most viable multi-tier option to support this architecture. The data points must be accessible from the REST API.
Which action meets these requirements for storing and retrieving location data?
A. Use Amazon Athena with Amazon S3.
B. Use Amazon API Gateway with AWS Lambda.
C. Use Amazon QuickSight with Amazon Redshift.
D. Use Amazon API Gateway with Amazon Kinesis Data Analytics.
Answer: D (LEAVE A REPLY)
Explanation/Reference: https://aws.amazon.com/kinesis/data-analytics/

NEW QUESTION: 202
You have set up Auto Scaling to automatically scale in. Consequently, you must decide which instances Auto Scaling should end first. What should you use to configure this?
A. An Elastic Load Balancer
B. A termination policy
C. An IAM role
D. Another scaling group
Answer: B (LEAVE A REPLY)
If you instruct Auto Scaling to automatically scale in, you must decide which instances Auto Scaling should terminate first. This can be configured through the use of a termination policy.
http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingBehavior.Insta nceTermination.html

NEW QUESTION: 203
A gaming company created a game leaderboard by using a Multi-AZ deployment of an Amazon RDS database. The number of users is growing, and the queries to get individual player rankings are getting slower over time. The company expects a surge in users for an upcoming version and wants to optimize the design for scalability and performance.
Which solution will meet these requirements?
A. Keep the leaderboard data in the RDS DB instance. Provision a Multi-AZ deployment of an Amazon ElastiCache for Redis cluster.
B. Add a read-only replica to the RDS DB instance. Add an RDS Proxy database proxy.

C. Stream the leaderboard data by using Amazon Kinesis Data Firehose with an Amazon S3 bucket as the destination. Query the S3 bucket by using Amazon Athena for the leaderboard.
D. Migrate the database to Amazon DynamoDB. Store the leader different tables. Use Apache HiveQL JOIN statements to build the leaderboard
Answer: D (LEAVE A REPLY)

NEW QUESTION: 204
A company has several applications running in an on-premises data center. The data center runs a mix of Windows and Linux VMs managed by VMware vCenter. A solution architect needs to create a plan to migrate the application to AWS. However, the solution architect discovers that the documentation for the applications is not up to date and that there are no complete infrastructure diagrams. The company's developers lack time to discuss their applications and current usage with the solutions architect.
What should the solutions architect do the gather the required information?
A. Use the AWS Migration Portfolio Assessment (MPA) tool to connect to each of the VMs to collect the configuration and utilization data.
B. Register the on-premises VMs with the AWS Migration Hub to collect configuration and utilization data.
C. Install the AWS Application Discovery Service on each of the VMs to collect the configuration and utilization data.
D. Deploy the AWS server migration service (AWS SMS) connector using the OVA image on the VMware cluster to collect configuration data from the VMs.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 205
A company is finalizing the architecture for its backup solution for applications running on AWS. All of the applications run on AWS and use at least two Availability Zones in each tier.
Company policy requires IT to durably store nightly backups f all its data in at least two locations: production and disaster recovery. The locations must be in different geographic regions. The company also needs the backup to be available to restore immediately at the production data center, and within 24 hours at the disaster recovery location. All backup processes must be fully automated.
What is the MOST cost-effective backup solution that will meet all requirements?
A. Back up all the data to Amazon S3 in the disaster recovery region. Use a lifecycle policy to move this data to Amazon Glacier in the production region immediately. Only the data is replicated; remove the data from the S3 bucket in the disaster recovery region.
B. Back up all the data to a large Amazon EBS volume attached to the backup media server in the production region. Run automated scripts to snapshot these volumes nightly, and copy these snapshots to the disaster recovery region.
C. Back up all the data to Amazon S3 in the production region. Set up cross-region replication of this S3 bucket to another region and set up a lifecycle policy in the second region to immediately move this data to Amazon Glacier.
D. Back up all the data to Amazon Glacier in the production region. Set up cross-region replication of this data to Amazon Glacier in the disaster recovery region. Set up a lifecycle policy to delete any data older than 60 days.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 206
A large company has many business units. Each business unit has multiple AWS accounts for different purposes. The CIO of the company sees that each business unit has data that would be useful to share with other parts of the company. In total, there are about 10 PB of data that needs to be shared with users in 1,000 AWS accounts. The data is proprietary, so some of it should only be available to users with specific job types. Some of the data is used for throughput of intensive workloads, such as simulations. The number of accounts changes frequently because of new initiatives, acquisitions, and divestitures.
A Solutions Architect has been asked to design a system that will allow for sharing data for use in AWS with all of the employees in the company.
Which approach will allow for secure data sharing in scalable way?
A. Store the data in a series of Amazon S3 buckets. Create an application running in Amazon EC2 that is integrated with the company's identity provider (IdP) that authenticates users and allows them to download or upload data through the application. The application uses the business unit and job type information in the IdP to control what users can upload and download through the application. The users can access the data through the application's API.
B. Store the data in a series of Amazon S3 buckets. Create an AWS STS token vending machine that is integrated with the company's identity provider (IdP). When a user logs in, have the token vending machine attach an IAM policy that assumes the role that limits the user's access and/or upload only the data the user is authorized to access. Users can get credentials by authenticating to the token vending machine's website or API and then use those credentials with an S3 client.
C. Store the data in a single Amazon S3 bucket. Create an IAM role for every combination of job type and business unit that allows to appropriate read/write access based on object prefixes in the S3 bucket. The roles should have trust policies that allow the business unit's AWS accounts to assume their roles. Use IAM in each business unit's AWS account to prevent them from assuming roles for a different job type. Users get credentials to access the data by using AssumeRole from their business unit's AWS account. Users can then use those credentials with an S3 client.
D. Store the data in a single Amazon S3 bucket. Write a bucket policy that uses conditions to grant read and write access where appropriate, based on each user's business unit and job type. Determine the business unit with the AWS account accessing the bucket and the job type with a prefix in the IAM user's name. Users can access data by using IAM credentials from their business unit's AWS account with an S3 client.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 207
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: B (LEAVE A REPLY)
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: 208
A company has a web application that allows users to upload short videos. The videos are stored on Amazon EBS volumes and analyzed by custom recognition software for categorization.
The website contains stat c content that has variable traffic with peaks in certain months. The architecture consists of Amazon EC2 instances running in an Auto Scaling group for the web application and EC2 instances running in an Auto Scaling group to process an Amazon SQS queue The company wants to re-architect the application to reduce operational overhead using AWS managed services where possible and remove dependencies on third-party software.
Which solution meets these requirements?
A. Host the web application in Amazon S3. Store the uploaded videos in Amazon S3. Use S3 event notifications to publish events to the SQS queue Process the SQS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
B. Use Amazon ECS containers for the web application and Spot Instances for the Auto Scaling group that processes the SQS queue. Replace the custom software with Amazon Recognition to categorize the videos.
C. Use AWS Elastic Beanstalk to launch EC2 instances in an Auto Scaling group for the web application and launch a worker environment to process the SQS queue Replace the custom software with Amazon Rekognition to categorize the videos.
D. Store the uploaded videos n Amazon EFS and mount the file system to the EC2 instances for Te web application. Process the SOS queue with an AWS Lambda function that calls the Amazon Rekognition API to categorize the videos.
Answer: (SHOW ANSWER)

NEW QUESTION: 209
A company has an internal AWS Elastic Beanstalk worker environment inside a VPC that must access an external payment gateway API available on an HTTPS endpoint the public internet Because of security policies, the payment gateway's Application team can grant access to only one public IP address Which architecture will

set up an Elastic Beanstalk environment to access the company's application without making multiple changes on the company's end?

A. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet with an internet gateway Associate an Elastic IP address to the internet gateway that can be whitelisted on the payment gateway application side

B. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet with an outbound route to a NAT gateway in a public subnet Associate an Elastic IP address to the NAT gateway that can be whitelisted on the payment gateway application side

C. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a public subnet Set the https_proxy and no_proxy application parameters to send non-VPC outbound HTTPS connections to an EC2 proxy server deployed in a public subnet Associate an Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gateway application side

D. Configure the Elastic Beanstalk application to place Amazon EC2 instances in a private subnet Set an https_proxy application parameter to send outbound HTTPS connections to an EC2 proxy server deployed in a public subnet Associate an Elastic IP address to the EC2 proxy host that can be whitelisted on the payment gateway application side

Answer: D (LEAVE A REPLY)

NEW QUESTION: 210

A solutions architect is designing a network for a new cloud deployment. Each account will need autonomy to modify route tables and make changes. Centralized and controlled egress internet connectivity is also needed. The cloud footprint is expected to grow to thousands ol AWS accounts.

Which architecture will meet these requirements?

A. A centralized transit VPC with a VPN connection to a standalone VPC in each account. Outbound internet traffic will be controlled by firewall appliances.

B. A shared services VPC to host central assets to include a fleet of firewalls wilh a route to the internet. Each spoke VPC will peer to the central VPC.

C. A centralized shared VPC with a subnet for each account. Outbound internet traffic will controlled through a fleet of proxy servers.

D. A shared transit gateway to which each VPC will be attached. Outbound internet access will route through a fleet of VPN-attached firewalls.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 211

A Solutions Architect is designing the storage layer for a data warehousing application. The data files are large, but they have statically placed metadata at the beginning of each file that describes the size and placement of the file's index. The data files are read in by a fleet of Amazon EC2 instances that store the index size, index location, and other category information about the data file in a database. That database is used by Amazon EMR to group files together for deeper analysis.

What would be the MOST cost-effective, high availability storage solution for this workflow?

A. Store the data files on Amazon EBS volumes and allow the EC2 fleet and EMR to mount and unmount the volumes where they are needed.
B. Store the data files in Amazon EFS mounted by the EC2 fleet and EMR nodes.
C. Store the content of the data files in Amazon DynamoDB tables with the metadata, index, and data as their own keys.
D. Store the data files in Amazon S3 and use Range GET for each file's metadata, then index the relevant data.
Answer: (SHOW ANSWER)

NEW QUESTION: 212
You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance tor a total of seven EC2 instances. The web, application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS Web (raffle gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.
Which of the following could be the root caused? (Choose 2 answers)
A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches
D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances
E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances
Answer: C,E (LEAVE A REPLY)
Explanation
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION: 213

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: (SHOW ANSWER)
Explanation
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: 214
A travel company built a web application that uses Amazon Simple Email Service (Amazon SES) to send email notifications to users. The company needs to enable logging to help troubleshoot email delivery issues.
The company also needs the ability to do searches that are based on recipient, subject, and time sent.
Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)
A. Use Amazon Athena to query the logs in Amazon CloudWatch for recipient, subject, and time sent.
B. Enable AWS CloudTrail logging. Specify an Amazon S3 bucket as the destination for the logs.
C. Create an Amazon SES configuration set with Amazon Kinesis Data Firehose as the destination.
Choose to send logs to an Amazon S3 bucket.
D. Create an Amazon CloudWatch log group. Configure Amazon SES to send logs to the log group.
E. Use Amazon Athena to query the logs in the Amazon S3 bucket for recipient, subject, and time sent.
Answer: B,E (LEAVE A REPLY)

NEW QUESTION: 215
A company is developing a new machine learning model solution in AWS. The models are developed as independent microservices that fetch about 1 GB of model data from Amazon S3 at startup and load the data into memory. Users access the models through an asynchronous API. Users can send a request or a batch of requests and specify where the results should be sent.
The company provides models to hundreds of users. The usage patterns for the models are irregular Some models could be unused for days or weeks. Other models could receive batches of thousands of requests at a time.
Which solution meets these requirements?
A. The requests from the API are sent to the model's Amazon Simple Queue Service (Amazon SQS) queue. Models are deployed as Amazon Elastic Container Service (Amazon ECS) services reading from the queue AWS App Mesh scales the instances of the ECS cluster based on the SQS queue size.
B. The requests from the API are sent to the models Amazon Simple Queue Service (Amazon SQS) queue. Models are deployed as AWS Lambda functions triggered by SQS events AWS Auto Scaling is enabled on Lambda to increase the number of vCPUs based on the SQS queue size.
C. The requests from the API are sent to the models Amazon Simple Queue Service (Amazon SQS) queue. Models are deployed as Amazon Elastic Container Service (Amazon ECS) services reading from the queue AWS Auto Scaling is enabled on Amazon ECS for both the cluster and copies of the service based on the queue size.
D. The requests from the API are sent to an Application Load Balancer (ALB). Models are deployed as AWS Lambda functions invoked by the ALB.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 216
A company is implementing a multi-account strategy; however, the Management team has expressed concerns that services like DNS may become overly complex. The company needs a solution that allows private DNS to be shared among virtual private clouds (VPCs) in different accounts. The company will have approximately 50 accounts in total.
What solution would create the LEAST complex DNS architecture and ensure that each VPC can resolve all AWS resources?
A. Create a shared services VPC in a central account, and create a VPC peering connection from the shared services VPC to each of the VPCs in the other accounts. Within Amazon Route 53, create a privately hosted zone in the shared services VPC and resource record sets for the domain and subdomains. Programmatically associate other VPCs with the hosted zone.
B. Create a VPC peering connection among the VPCs in all accounts. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "true" for each VPC. Create an Amazon Route 53 private zone for each VPC. Create resource record sets for the domain and subdomains.
Programmatically associate the hosted zones in each VPC with the other VPCs.
C. Create a shared services VPC in a central account. Create a VPC peering connection from the VPCs in other accounts to the shared services VPC. Create an Amazon Route 53 privately hosted zone in the shared services

VPC with resource record sets for the domain and subdomains. Allow UDP and TCP port 53 over the VPC peering connections.
D. Set the VPC attributes enableDnsHostnames and enableDnsSupport to "false" in every VPC. Create an AWS Direct Connect connection with a private virtual interface. Allow UDP and TCP port 53 over the virtual interface. Use the on-premises DNS servers to resolve the IP addresses in each VPC on AWS.
Answer: A (LEAVE A REPLY)
Explanation
https://aws.amazon.com/blogs/networking-and-content-delivery/centralized-dns-management-of-hybrid-cloud-w

NEW QUESTION: 217
A company is running multiple applications on Amazon EC2. Each application is deployed and managed by multiple business units. All applications are deployed on a single AWS account but on different virtual private clouds (VPCs). The company uses a separate VPC in the same account for test and development purposes. Production applications suffered multiple outages when users accidentally terminated and modified resources that belonged to another business unit. A Solutions Architect has been asked to improve the availability of the company applications while allowing the Developers access to the resources they need.
Which option meets the requirements with the LEAST disruption?
A. Create an AWS account for each business unit. Move each business unit's instances to its own account and set up a federation to allow users to access their business unit's account.
B. Set up a federation to allow users to use their corporate credentials, and lock the users down to their own VPC. Use a network ACL to block each VPC from accessing other VPCs.
C. Implement a tagging policy based on business units. Create an IAM policy so that each user can terminate instances belonging to their own business units only.
D. Set up role-based access for each user and provide limited permissions based on individual roles and the services for which each user is responsible.
Answer: C (LEAVE A REPLY)
Explanation
Principal - Control what the person making the request (the principal) is allowed to do based on the tags that are attached to that person's IAM user or role. To do this, use the aws:PrincipalTag/key-name condition key to specify what tags must be attached to the IAM user or role before the request is allowed.
https://docs.aws.amazon.com/IAM/latest/UserGuide/access_iam-tags.html

NEW QUESTION: 218
A company runs a proprietary stateless ETL application on an Amazon EC2 Linux instance. The application is a Linux binary, and the source code cannot be modified. The application is single-threaded, uses 2 GB of RAM. and is highly CPU intensive The application is scheduled to run every 4 hours and runs for up to 20 minutes A solutions architect wants to revise the architecture for the solution.
Which strategy should the solutions architect use?
A. Use AWS Lambda to run the application. Use Amazon CloudWatch Logs to invoke the Lambda function every 4 hours.

B. Use AWS Batch to run the application. Use an AWS Step Functions state machine to invoke the AWS Batch job every 4 hours.
C. Use AWS Fargate to run the application. Use Amazon EventBridge (Amazon CloudWatch Events) to invoke the Fargate task every 4 hours.
D. Use Amazon EC2 Spot Instances to run the application. Use AWS CodeDeploy to deploy and run the application every 4 hours.
Answer: C (LEAVE A REPLY)
Explanation
step function could run a scheduled task when triggered by eventbrige, but why would you add that layer of complexity just to run aws batch when you could directly invoke it through eventbridge. The link provided - https://aws.amazon.com/pt/blogs/compute/orchestrating-high-performance-computing-with-aws-step-functions-a makes sense only for HPC, this is a single instance that needs to be run

NEW QUESTION: 219
A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS lor PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.
What changes to the current architecture will reduce operational overhead and support the product release?
A. Create an EC2 Auto Scaling group behind an Application Load Balancer. Create additional read replicas for the DB instance. Create Amazon Kinesis data streams and configure the application services lo use the data streams. Store and serve static content directly from Amazon S3.
B. Create an EC2 Auto Scaling group behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.
C. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
D. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
Answer: (SHOW ANSWER)
Explanation
Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

NEW QUESTION: 220

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)
Explanation
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 VPC ensures 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: 221
A company hosts historical weather records in Amazon S3. The records are downloaded from the company's website by a way of a URL that resolves to a domain name. Users all over the world access this content through subscriptions. A third-party provider hosts the company's root domain name, but the company recently migrated some of its services to Amazon Route 53. The company wants to consolidate contracts, reduce latency for users, and reduce costs related to serving the application to subscribers.
Which solution meets these requirements?
A. Create a web distribution on Amazon CloudFront to serve the S3 content for the application. Create a CNAME record in a Route 53 hosted zone that points to the CloudFront distribution, resolving to the application's URL domain name.

B. Create an A record in a Route 53 hosted zone for the application. Create a Route 53 traffic policy for the web application, and configure a geoproximity rule. Configure health checks to check the health of the endpoint and route DNS queries to other endpoints if an endpoint is unhealthy.
C. Create an A record in a Route 53 hosted zone for the application. Create a Route 53 traffic policy for the web application, and configure a geolocation rule. Configure health checks to check the health of the endpoint and route DNS queries to other endpoints if an endpoint is unhealthy.
D. Create a web distribution on Amazon CloudFront to serve the S3 content for the application. Create an ALIAS record in the Amazon Route 53 hosted zone that points to the CloudFront distribution, resolving to the application's URL domain name.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 222
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: 223

A company has used infrastructure as code (laC) to provision a set of two Amazon EC2 instances. The instances have remained the same tor several years.
The company's business has grown rapidly in the past few months. In response, the company's operations team has implemented an Auto Scaling group to manage the sudden increases in traffic Company policy requires a monthly installation of security updates on all operating systems that are running.
The most recent security update required a reboot. As a result the Auto Scaling group terminated the instances and replaced them with new, unpatched instances.
Which combination of steps should a sol-tons architect recommend to avoid a recurrence of this issue? (Select TWO)
A. Create an Elastic Load Balancer in front of the Auto Scaling group Configure termination protection on the instances.
B. Modify the Auto Scaling group by setting the Update policy to target the oldest launch configuration for replacement.
C. Create an Elastic Load Balancer in front of the Auto Scaling group Configure monitoring to ensure that target group health checks return healthy after the Auto Scaling group replaces the terminated instances
D. Create a new Auto Scaling group before the next patch maintenance During the maintenance window patch both groups and reboot the instances.
E. Create automation scripts to patch an AMI. update the launch configuration, and invoke an Auto Scaling instance refresh.
Answer: B,C (LEAVE A REPLY)

NEW QUESTION: 224
How can you check the operational validity of your AWS CloudFormation template?
A. To check the operational validity, you need to attempt to create the stack.
B. There is no way to check the operational validity of your AWS CloudFormation template.
C. To check the operational validity, you need a sandbox or test area for AWS CloudFormation stacks.
D. To check the operational validity, you need to use the aws cloudformation validate-template command.
Answer: A (LEAVE A REPLY)
In AWS CloudFormation, to check the operational validity, you need to attempt to create the stack. There is no sandbox or test area for AWS CloudFormation stacks, so you are charged for the resources you create during testing.
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-validate- template.html

NEW QUESTION: 225
A development team has created a new flight tracker application that provides near-real-time data to users. The application has a front end that consists of an Application Load Balancer (ALB) in front of two large Amazon EC2 instances in a single Availability Zone. Data is stored in a single Amazon RDS MySQL DB instance. An Amazon Route 53 DNS record points to the ALB.
Management wants the development team to improve the solution to achieve maximum reliability with the least amount of operational overhead.
Which set of actions should the team take?

A. Configure the DB instance as Multi-AZ. Deploy the application to two additional EC2 instances in different Availability Zones behind an ALB.
B. Create RDS MySQL read replicas. Deploy the application to multiple AWS Regions. Use a Route 53 latency-based routing policy to route to the application.
C. Replace the DB instance with Amazon DynamoDB global tables. Deploy the application in multiple AWS Regions. Use a Route 53 latency-based routing policy to route to the application.
D. Replace the DB instance with Amazon Aurora with Aurora Replicas. Deploy the application to mulliple smaller EC2 instances across multiple Availability Zones in an Auto Scaling group behind an ALB.
Answer: (SHOW ANSWER)

NEW QUESTION: 226
A company has a policy that all Amazon EC2 instances that are running a database must exist within the same subnets in a shared VPC Administrators must follow security compliance requirements and are not allowed to directly log in to the shared account All company accounts are members of the same organization in AWS Organizations. The number of accounts will rapidly increase as the company grows.
A. A solutions architect uses AWS Resource Access Manager to create a resource share in the shared account. What is the MOST operationally efficient configuration to meet these requirements?
B. Add all subnets within the VPC to the resource share Add the organization as a principal
C. Add the VPC to the resource share Add the organization as a principal
D. Add all subnets within the VPC to the resource share Add the account IDs as principals
E. Add the VPC to the resource share Add the account IDs as principals
Answer: (SHOW ANSWER)

NEW QUESTION: 227
A solutions architect is designing storage for a high performance computing (HPC) environment based on Amazon Linux. The workload stores and processes a large amount of engineering drawings that require shared storage and heavy computing.
Which storage option would be the optimal solution?
A. Amazon Elastic File System (Amazon EFS)
B. Amazon FSx for Lustre
C. Amazon EC2 instance store

D. Amazon EBS Provisioned IOPS SSD (io1)

Answer: B (LEAVE A REPLY)

Explanation/Reference: https://d1.awsstatic.com/whitepapers/AWS%20Partner%20Network_HPC%20Storage%20Options_2019_FINAL.pdf (p.8)

NEW QUESTION: 228

A company has a website that enables users to upload videos Company policy states the uploaded videos must be analyzed for restricted content An uploaded video is placed in Amazon S3, and a message is pushed to an Amazon SQS queue with the video's location A backend application pulls this location from Amazon SQS and analyzes the video The video analysis is compute-intensive and occurs sporadically during the day The website scales with demand The video analysis application runs on a fixed number of instances Peak demand occurs during the holidays, so the company must add instances to the application during this time All instances used are currently on-demand Amazon EC2 T2 instances The company wants to reduce the cost of the current solution.

Which of the following solutions is MOST cost-effective?

A. Keep the website on T2 instances Determine the minimum number of website instances required during off-peak times and use Spot Instances to cover them while using Reserved Instances to covet peak demand Use Amazon EC2 R4 and Amazon EC2 R5 Reserved Instances in an Auto Scaling group for the video analysis application

B. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 C4 instances Determine the minimum number of website instances required during off-peak times and use On-Demand instances to cover them while using Spot capacity to cover peak demand Use Spot Fleet for the video analysis application comprised of C4 and Amazon EC2 C5 instances

C. Migrate the website to AWS Elastic Beanstalk and Amazon EC2 R4 instances Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand Use Spot Fleet for the video analysis application comprised of R4 and Amazon EC2 R5 instances.

D. Keep the website on 12 instances Determine the minimum number of website instances required during off-peak times and use Reserved Instances to cover them while using On-Demand Instances to cover peak demand Use Spot Fleet for the video analysis application comprised of Amazon EC2 C4 and Amazon EC2 C5 Spot Instances

Answer: D (LEAVE A REPLY)

NEW QUESTION: 229

A company's interactive web application uses an Amazon CloudFront distribution to serve images from an Amazon S3 bucket. Occasionally, third-party tools ingest corrupted images into the S3 bucket. This image corruption causes a poor user experience in the application later. The company has successfully implemented and tested Python logic to detect corrupt images.

A solutions architect must recommend a solution to integrate the detection logic with minimal latency between the ingestion and serving.

Which solution will meet these requirements?

A. Use a Lambda@Edge function that is invoked by a viewer-response event.

B. Use a Lambda@Edge function that is invoked by an origin-response event.
C. Use an S3 event notification that invokes an AWS Lambda function.
D. Use an S3 event notification that invokes an AWS Step Functions state machine.
Answer: (SHOW ANSWER)
This solution will allow the detection logic to be run as soon as the image is uploaded to the S3 bucket, before it is served to users via the CloudFront distribution. This way, the detection logic can quickly identify any corrupted images and prevent them from being served to users, minimizing latency between ingestion and serving.
Reference: AWS Lambda@Edge documentation:
https://docs.aws.amazon.com/lambda/latest/dg/lambda-edge.html You can use Lambda@Edge to run your code in response to CloudFront events, such as a viewer request, an origin request, a response, or an error.

NEW QUESTION: 230
What is the maximum number of data points for an HTTP data request that a user can include in PutMetricRequest in the CloudWatch?
A. 30
B. 50
C. 10
D. 20
Answer: D (LEAVE A REPLY)
The size of a PutMetricData request of CloudWatch is limited to 8KB for the HTTP GET requests and 40KB for the HTTP POST requests. The user can include a maximum of 20 data points in one PutMetricData request.
http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/cloudwatch_concepts.ht ml

NEW QUESTION: 231
In the context of Amazon ElastiCache CLI, which of the following commands can you use to view all ElastiCache instance events for the past 24 hours?
A. elasticache-events --duration 24
B. elasticache-events --duration 1440
C. elasticache-describe-events --duration 24
D. elasticache describe-events --source-type cache-cluster --duration 1440
Answer: D (LEAVE A REPLY)
Explanation
In Amazon ElastiCache, the code "aws elasticache describe-events --source-type cache-cluster -- duration 1440" is used to list the cache-cluster events for the past 24 hours (1440 minutes).
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/ECEvents.Viewing.html

NEW QUESTION: 232
A media company has a static web application that is generated programmatically. The company has a build pipeline that generates HTML content that is uploaded to an Amazon S3 bucket served by Amazon CloudFront. The build pipeline runs inside a Build Account. The S3 bucket and CloudFront distribution are in a Distribution Account. The build pipeline uploads the files to Amazon S3 using an IAM role in the Build Account. The S3

bucket has a bucket policy that only allows CloudFront to read objects using an origin access identity (OAI). During testing, all attempts to access the application using the CloudFront URL result in an HTTP 403 Access Denied response.
What should a solutions architect suggest to the company to allow access the objects in Amazon S3 through CloudFront?
A. Modify the S3 upload process in the Build Account to add the bucket-owner-full-control ACL to the objects at upload.
B. Create a new cross-account IAM role in the Distribution Account with write access to the S3 bucket. Modify the build pipeline to assume this role to upload the files to the Distribution Account.
C. Modify the S3 upload process in the Build Account to set the object owner to the Distribution Account.
D. Create a new IAM role in the Distribution Account with read access to the S3 bucket. Configure CloudFront to use this new role as its OAI. Modify the build pipeline to assume this role when uploading files from the
Answer: A (LEAVE A REPLY)
Explanation
https://aws.amazon.com/fr/premiumsupport/knowledge-center/s3-bucket-owner-access/

NEW QUESTION: 233
A company is selling up an application to use an Amazon RDS MySQL DB instance. The database must be architected for high availability across Availability Zones and AWS Regions with minimal downtime.
How should a solutions architect meet this requirement?
A. Set up an RDS MySQL Single-AZ DB instance. Copy automated snapshots to at least one other Region.
B. Set up an RDS MySQL Single-AZ DB instance. Configure a read replica in a different Region.
C. Set up an RDS MySQL Multi-AZ DB instance. Configure a read replica in a different Region.
D. Set up an RDS MySQL Multi-AZ DB instance. Configure an appropriate backup window.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 234
A security engineer determined that an existing application retrieves credentials to an Amazon RDS for MySQL database from an er 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 e 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 SecureString parameter type using AWS Systems Manager Parameter Store. Specify an AWS AppSync DataSource resource to automatically rotate the database password.
D. Generate the database password as a secret resource using AWS Secrets Manager. Create an AWS Lambda function resource to rotate the database password. Create an Amazon EventBridge scheduled rule resource to trigger the Lambda function password rotation every 90 days.
Answer: (SHOW ANSWER)

NEW QUESTION: 235
How can a user list the IAM Role configured as a part of the launch config?
A. as-describe-launch-configs -iam-profile
B. as-describe-launch-configs -show-long
C. as-describe-launch-configs -iam-role
D. as-describe-launch-configs -role
Answer: (SHOW ANSWER)
As-describe-launch-configs describes all the launch config parameters created by the AWS account in the specified region. Generally it returns values, such as Launch Config name, Instance Type and AMI ID. If the user wants additional parameters, such as the IAM Profile used in the config , he has to run command:
as-describe-launch-configs --show-long

NEW QUESTION: 236
A user is trying to create a PIOPS EBS volume with 3 GB size and 90 IOPS. Will AWS create the volume?
A. No, since the PIOPS and EBS size ratio is less than 30
B. Yes, since the ratio between EBS and IOPS is less than 30
C. No, the EBS size is less than 4GB
D. Yes, since PIOPS is higher than 100
Answer: C (LEAVE A REPLY)
Explanation
A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html#EBSVolumeTypes_piops

NEW QUESTION: 237
A solutions architect is working with a company that is extremely sensitive to its IT costs and wishes to implement controls that will result in a predictable AWS spend each month Which combination ot steps can help the company control and monitor its monthly AWS usage to achieve a cost that is as close as possible to the target amount? (Select THREE.)
A. Define 'workload' as a cost allocation tag in the AWS Billing and Cost Management console
B. Contact AWS Support and ask that they apply limits to the account so that users are not able to launch more than a certain number of instance types
C. Implement an IAM policy that requires users to specify a 'workload' tag for cost allocation when launching Amazon EC2 instances

D. Place conditions in the users' IAM policies that limit the number of instances they are able to launch
E. Set up AWS Budgets to alert and notify when a given workload is expected to exceed a defined cost
F. Purchase all upfront Reserved Instances that cover 100% of the account's expected Amazon EC2 usage
Answer: A,C,E (LEAVE A REPLY)

NEW QUESTION: 238
A Solutions Architect must design a highly available, stateless, REST service. The service will require multiple persistent storage layers for service object meta information and the delivery of content. Each request needs to be authenticated and securely processed. There is a requirement to keep costs as low as possible?
How can these requirements be met?
A. Use AWS Fargate to host a container that runs a self-contained REST service. Set up an Amazon ECS service that is fronted by an Application Load Balancer (ALB). Use a custom authenticator to control access to the API. Store request meta information in Amazon DynamoDB with Auto Scaling and static content in a secured S3 bucket. Make secure signed requests for Amazon S3 objects and proxy the data through the REST service interface.
B. Use AWS Fargate to host a container that runs a self-contained REST service. Set up an ECS service that is fronted by a cross-zone ALB. Use an Amazon Cognito user pool to control access to the API.
Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.
C. Set up Amazon API Gateway and create the required API resources and methods. Use an Amazon Cognito user pool to control access to the API. Configure the methods to use AWS Lambda proxy integrations, and process each resource with a unique AWS Lambda function.
Store request meta information in DynamoDB with Auto Scaling and static content in a secured S3 bucket. Generate presigned URLs when returning references to content stored in Amazon S3.
D. Set up Amazon API Gateway and create the required API resources and methods. Use an Amazon API Gateway custom authorizer to control access to the API. Configure the methods to use AWS Lambda custom integrations, and process each resource with a unique Lambda function. Store request meta information in an Amazon ElastiCache Multi-AZ cluster and static content in a secured S3 bucket.
Generate presigned URLs when returning references to content stored in Amazon S3.
Answer: C (LEAVE A REPLY)
The Cognito can be used for authentication. API Gateway for Stateless REST service.
A: uses custom authenticator. Self Contained REST service!!!
B: Self Contained REST service!!!
D: Elasticache is not persistent.

NEW QUESTION: 239
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. Convert the organization's root SCPs from deny list SCPs to allow list SCPs to allow the required services only. Temporally apply an SCP to the organization's root that allows AWS Config actions for principals only in 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 organization's root SCP to the Production OU. Move the new account to the Production OU when adjustments to AWS Config are complete
C. 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
D. Create a temporary OU named Onboarding for the new account. Apply an SCP to the Onboardmg OU to allow AWS Config actions Move the new account to the Production OU when adjustments to AWS Config are complete
Answer: D (LEAVE A REPLY)

NEW QUESTION: 240
Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fleet of spot EC2 instances. Files submitted by your premium customers must be transformed with the highest priority. How should you implement such a system?
A. Use a single SQS queue. Each message contains the priority level. Transformation instances poll high-priority messages first.
B. Use Route 53 latency based-routing to send high priority tasks to the closest transformation instances.
C. Use a DynamoDB table with an attribute defining the priority level. Transformation instances will scan the table for tasks, sorting the results by priority level.
D. Use two SQS queues, one for high priority messages, the other for default priority. Transformation instances first poll the high priority queue; if there is no message, they poll the default priority queue.
Answer: D (LEAVE A REPLY)

NEW QUESTION: 241
An organization is trying to setup a VPC with Auto Scaling. Which configuration steps below is not required to setup AWS VPC with Auto Scaling?
A. Configure the Auto Scaling group with the VPC ID in which instances will be launched.
B. Configure the Auto Scaling Launch configuration with multiple subnets of the VPC to enable the Multi AZ feature.
C. Configure the Auto Scaling Launch configuration which does not allow assigning a public IP to instances.
D. Configure the Auto Scaling Launch configuration with the VPC security group.

Answer: B (LEAVE A REPLY)

Explanation

The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an Auto Scaling group. Before creating the Auto Scaling group it is recommended that the user creates the Launch configuration. Since it is a VPC, it is recommended to select the parameter which does not allow assigning a public IP to the instances.

The user should also set the VPC security group with the Launch configuration and select the subnets where the instances will be launched in the AutoScaling group. The HA will be provided as the subnets may be a part of separate AZs.

http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/autoscalingsubnets.html

NEW QUESTION: 242

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apacheweb server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database. Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis. Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events? (Select THREE.)

A. Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.

B. Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis.

C. Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.

D. Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs

E. Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.

F. Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Answer: (SHOW ANSWER)

NEW QUESTION: 243
A company is running a web application in a VPC. The web application runs on a group of Amazon EC2 instances behind an Application Load Balancer (ALB). The ALB is using AWS WAF.
An external customer needs to connect to the web application. The company must provide IP addresses to all external customers.
Which solution will meet these requirements with the LEAST operational overhead?
A. Replace the ALB with a Network Load Balancer (NLB). Assign an Elastic IP address to the NLB.
B. Allocate an Elastic IP address. Assign the Elastic IP address to the ALProvide the Elastic IP address to the customer.
C. Create an AWS Global Accelerator standard accelerator. Specify the ALB as the accelerator's endpoint. Provide the accelerator's IP addresses to the customer.
D. Configure an Amazon CloudFront distribution. Set the ALB as the origin. Ping the distribution's DNS name to determine the distribution's public IP address. Provide the IP address to the customer.
Answer: C (LEAVE A REPLY)
Explanation
https://docs.aws.amazon.com/global-accelerator/latest/dg/about-accelerators.alb-accelerator.html Option A is wrong. AWS WAF does not support associating with NLB.
https://docs.aws.amazon.com/waf/latest/developerguide/waf-chapter.html Option B is wrong. An ALB does not support an Elastic IP address. https://aws.amazon.com/elasticloadbalancing/features/

NEW QUESTION: 244
A company is migrating a three-tier application to AWS. The application requires a MySQL database. In the past, the application users reported poor application performance when creating new entries. These performance issues were caused by users generating different real-time reports from the application during working hours.
Which solution will improve the performance of the application when it is moved to AWS?
A. Create an Amazon Aurora MySQL Multi-AZ DB cluster. Configure the application to use the backup instance of the cluster as an endpoint for the reports.
B. Create an Amazon Aurora MySQL Multi-AZ DB cluster with multiple read replicas. Configure the application reader endpoint for reports.
C. Import the data into an Amazon DynamoDB table with provisioned capacity. Refactor the application to use DynamoDB for reports.
D. Create the database on a compute optimized Amazon EC2 instance. Ensure compute resources exceed the on-premises database.
Answer: B (LEAVE A REPLY)

NEW QUESTION: 245
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 the networking account Peer VPCs from the accounts that host the S3 buckets with the VPC in the network account

Answer: A,C (LEAVE A REPLY)

Explanation

https://aws.amazon.com/premiumsupport/knowledge-center/s3-bucket-access-direct-connect/

NEW QUESTION: 246

Example Corp. has an on-premises data center and a VPC named VPC A in the Example Corp. AWS account. The on-premises network connects to VPC A through an AWS Site-To-Site VPN. The on-premises servers can properly access VPC A. Example Corp. just acquired AnyCompany, which has a VPC named VPC B.

There is no IP address overlap among these networks. Example Corp. has peered VPC A and VPC B.

Example Corp. wants to connect from its on-premise servers to VPC B. Example Corp. has properly set up the network ACL and security groups.

Which solution will meet this requirement with the LEAST operational effort?

A. Create a transit gateway. Attach the Site-to-Site VPN, VPC A, and VPC B to the transit gateway.
Update the transit gateway route tables for all networks to add IP range routes for all other networks.

B. Create a transit gateway. Create a Site-to-Site VPN connection between the on-premises network and VPC B. and connect the VPN connection to the transit gateway. Add a route to direct traffic to the peered VPCs, and add an authorization rule to give clients access to the VPCs A and B.

C. Update the route tables for the Site-to-Site VPN and both VPCs for all three networks. Configure BGP propagation for all three networks. Wait for up to 5 minutes for BGP propagation to finish.

D. Modify the Site-to-Site VPN's virtual private gateway definition to include VPC A and VPC B. Split the two routers of the virtual private getaway between the two VPCs.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 247

True or False: In Amazon ElastiCache replication groups of Redis, for performance tuning reasons, you can change the roles of the cache nodes within the replication group, with the primary and one of the replicas exchanging roles.

A. True, however, you get lower performance.

B. FALSE

C. TRUE

D. False, you must recreate the replication group to improve performance tuning.

Answer: C (LEAVE A REPLY)
In Amazon ElastiCache, a replication group is a collection of Redis Cache Clusters, with one primary read-write cluster and up to five secondary, read-only clusters, which are called read replicas. You can change the roles of the cache clusters within the replication group, with the primary cluster and one of the replicas exchanging roles. You might decide to do this for performance tuning reasons.
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/Replication.Redis.Groups.ht ml

NEW QUESTION: 248
A company has a mobile app with users in Europe. When the app is used, it downloads a configuration file that is device and app version-specific. The company has the following architecture:
* Configuration files are stored in Amazon S3 in the eu-west-1 Region and served to the users using Amazon CloudFront.
* Lambda@Edge is used to extract the device and version information from the app requests. It then updates the requests to load the correct configuration.
The company uses the configuration file load time as a key performance metric, and targets a response time of 100 ms or less. The app recently launched in the ap-southeast-2 Region, and the latency for requests from users in Australia is significantly above the 100 ms target.
A solutions architect needs to recommend a solution.
Which solution will reduce latency for users in Australia?
A. Create an S3 bucket in the ap-southeast-2 Region. Use cross-Region replication to synchronize from the bucket in the eu-west-1 Region Modify Lambda@Edge to access Amazon S3 in the Region that is closest to the user.
B. Configure S3 Transfer Acceleration on the bucket. Add the Transfer Acceleration Edge endpoints for Australia and Europe as CloudFront origins. Modify Lambda@Edge to update the origin of the request to be the Transfer Acceleration endpoint in the Region that is closest to the user.
C. Configure S3 Transfer Acceleration on the bucket. Modify Lambda@Edge to access Amazon S3 using the Transfer Acceleration endpoint in the Region that is closest to the user.
D. Create an S3 bucket in the ap-southeast-2 Region. Use cross-Region replication to synchronize from the bucket in the eu-west-1 Region. Create an Amazon Route 53 hosted zone with latency-based routing configured for both buckets. Modify Lambda@Edge to update the origin of the request to be the Route 53 hosted zone that is closest to the user.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 249
A company provides AWS solutions to its users with AWS CloudFormation templates. Users launch the templates in their accounts to have different solutions provisioned for them. The users want to improve the deployment strategy for solutions while retaining the ability to do the following:
* Add their own features to a solution for their specific deployments.
* Run unit tests on their changes.
* Turn features on and off for their deployments.
* Automatically update with code changes.

* Run security scanning tools for their deployments.
Which strategies should the solutions architect use to meet the requirements?
A. Allow users to download solution code artifacts in their Amazon S3 buckets. Use Amazon S3 and AWS CodePipeline for the CI/CD pipelines. Use CloudFormation StackSets for different solution features and to turn features on and off. Use AWS Lambda to run unit tests and security scans, and CloudFormation for deploying and updating a solution with changes.
B. Allow users to download solution code artifacts. Use AWS CodeCommit and AWS CodePipeline for the CI/CD pipeline. Use AWS Amplify plugins for different solution features and user prompts to turn features on and off. Use AWS Lambda to run unit tests and security scans, and AWS CodeBuild for deploying and updating a solution with changes
C. Allow users to download solution code artifacts. Use AWS CodeCommit and AWS CodePipeline for the CI/CD pipeline. Use the AWS Cloud Development Kit constructs for different solution features, and use the manifest file to turn features on and off. Use AWS CodeBuild to run unit tests and security and for deploying and updating a solution with changes.
D. Allow users to download solution code as Docker images. Use AWS CodeBuild and AWS CodePipeline for the CI/CD pipeline. Use Docker images for different solution features and the AWS CLI to turn features on and off. Use AWS CodeDeploy to run unit tests and security scans, and for deploying and updating a solution with changes.
Answer: (SHOW ANSWER)

NEW QUESTION: 250
A company is providing weather data over a REST-based API to several customers. The API is hosted by Amazon API Gateway and is integrated with different AWS Lambda functions for each API operation. The company uses Amazon Route 53 for DNS and has created a resource record of weather.example.com. The company stores data for the API in Amazon DynamoDB tables. The company needs a solution that will give the API the ability to fail over to a different AWS Region.
Which solution will meet these requirements?
A. Deploy a new set of Lambda functions in a new Region. Update the API Gateway API to use an edge-optimized API endpoint with Lambda functions from both Regions as targets. Convert the DynamoDB tables to global tables.
B. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.
C. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a failover record. Enable target health monitoring. Convert the DynamoDB tables to global tables.
D. Deploy a new API Gateway API in a new Region. Change the Lambda functions to global functions. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.
Answer: C (LEAVE A REPLY)
https://docs.aws.amazon.com/apigateway/latest/developerguide/dns-failover.html

NEW QUESTION: 251
You create a VPN connection, and your VPN device supports Border Gateway Protocol (BGP). Which of the following should be specified to configure the VPN connection?
A. Classless routing
B. Classfull routing
C. Dynamic routing
D. Static routing
Answer: C (LEAVE A REPLY)
If you create a VPN connection, you must specify the type of routing that you plan to use, which will depend upon on the make and model of your VPN devices. If your VPN device supports Border Gateway Protocol (BGP), you need to specify dynamic routing when you configure your VPN connection. If your device does not support BGP, you should specify static routing.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

NEW QUESTION: 252
A company is migrating an application to AWS. It wants to use fully managed services as much as possible during the migration. The company needs to store large, important documents within the application with the following requirements:
* The data must be highly durable and available.
* The data must always be encrypted at rest and in transit.
* The encryption key must be managed by the company and rotated periodically.
Which of the following solutions should the Solutions Architect recommend?
A. Deploy the storage gateway to AWS in file gateway mode. Use Amazon EBS volume encryption using an AWS KMS key to encrypt the storage gateway volumes.
B. Use Amazon S3 with a bucket policy to enforce HTTPS for connections to the bucket and to enforce server-side encryption and AWS KMS for object encryption.
C. Use Amazon DynamoDB with SSL to connect to DynamoDB. Use an AWS KMS key to encrypt DynamoDB objects at rest.
D. Deploy instances with Amazon EBS volumes attached to store this data. Use EBS volume encryption using an AWS KMS key to encrypt the data.
Answer: B (LEAVE A REPLY)
https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-your-amazon-s3-data/

NEW QUESTION: 253
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: A (LEAVE A REPLY)
Explanation
https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/
https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/ The solutions architect should recommend implementing retry logic with exponential backoff and irregular variation in the client application. This will help reduce the number of errors by retrying failed requests with increasingly longer waits between attempts. Additionally, it is important to ensure that errors are caught and handled with descriptive error messages so that customers are not presented with an ambiguous error.

NEW QUESTION: 254
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 (NIB) foe scalability.
Up to 10 business unit VPCs will need to be connected to the shared VPC Some of 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 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 service console
B. 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
C. 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

D. 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

Answer: A (LEAVE A REPLY)

NEW QUESTION: 255

A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.

What AWS architecture would you recommend?

A. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket tor each customer with a Bucket Policy that permits access only to that one customer.

B. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshold. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.

C. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.

D. ASK their customers to use an S3 client instead of an FTP client. Create a single S3 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.

Answer: D (LEAVE A REPLY)