Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
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: D

Explanation:
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 push notification to all devices associated with an identity to notify them that new data is available.
Reference: http://aws.amazon.com/cognito/faqs/

NEW QUESTION 2
An organization is planning to extend their data center by connecting their DC with the AWS VPC using the VPN gateway. The organization is setting up a dynamically routed VPN connection. Which of the below mentioned answers is not required to setup this configuration?

A. The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha.
B. Elastic IP ranges that the organization wants to advertise over the VPN connection to the VPC.
C. Internet-routable IP address (static) of the customer gateway's external interface.
D. Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gatewa

Answer: B

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. The organization wants to extend their network into the cloud and also directly access the internet from their AWS VPC. Thus, the organization should setup a Virtual Private Cloud (VPC) with a public subnet and a private subnet, and a virtual private gateway to enable communication with their data center network over an IPsec VPN tunnel. To setup this configuration the organization needs to use the Amazon VPC with a VPN connection. The organization network administrator must designate a physical appliance as a customer gateway and configure it. The organization would need the below mentioned information to setup this configuration:
The type of customer gateway, such as Cisco ASA, Juniper J-Series, Juniper SSG, Yamaha Internet-routable IP address (static) of the customer gateway's external interface
Border Gateway Protocol (BGP) Autonomous System Number (ASN) of the customer gateway, if the organization is creating a dynamically routed VPN connection.
Internal network IP ranges that the user wants to advertise over the VPN connection to the VPC. Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

NEW QUESTION 3
An organization is setting a website on the AWS VPC. The organization has blocked a few IPs to avoid a D-DOS attack. How can the organization configure that a request from the above mentioned IPs does not access the application instances?

A. Create an IAM policy for VPC which has a condition to disallow traffic from that IP address.
B. Configure a security group at the subnet level which denies traffic from the selected IP.
C. Configure the security group with the EC2 instance which denies access from that IP address.
D. Configure an ACL at the subnet which denies the traffic from that IP addres

Answer: D

Explanation:
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. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. Security group works at the instance level while ACL works at the subnet level. ACL allows both allow and deny rules.
Thus, when the user wants to reject traffic from the selected IPs it is recommended to use ACL with subnets.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLs.html

NEW QUESTION 4
In which step of using AWS Direct Connect should the user determine the required port speed?

A. Complete the Cross Connect
B. Verify Your Virtual Interface
C. Download Router Configuration
D. Submit AWS Direct Connect Connection Request

Answer: D

Explanation:
To submit an AWS Direct Connect connection request, you need to provide the following information: Your contact information.
The AWS Direct Connect Location to connect to.
Details of AWS Direct Connect partner if you use the AWS Partner Network (APN) service. The port speed you require, either 1 Gbps or 10 Gbps.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.html#ConnectionRequest

NEW QUESTION 5
In Amazon IAM, what is the maximum length for a role name?

A. 128 characters
B. 512 characters
C. 64 characters
D. 256 characters

Answer: C

Explanation:
In Amazon IAM, the maximum length for a role name is 64 characters.
Reference: http://docs.aws.amazon.com/IANI/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 6
An organization is creating a VPC for their application hosting. The organization has created two private subnets in the same AZ and created one subnet in a separate zone. The organization wants to make a
HA system with the internal ELB. Which of these statements is true with respect to an internal ELB in this scenario?

A. ELB can support only one subnet in each availability zone.
B. ELB does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.
C. If the user is creating an internal ELB, he should use only private subnets.
D. ELB can support all the subnets irrespective of their zone

Answer: A

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 ELB, and EC2 instances.
There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App sewers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.
The Internal ELB supports only one subnet in each AZ and asks the user to select a subnet while configuring internal ELB.
Reference: http://docs.aws.amazon.com/EIasticLoadBaIancing/latest/DeveIoperGuide/USVPC_creating_basic_Ib.html

NEW QUESTION 7
A user is configuring MySQL RDS with PIOPS. What should be the minimum size of DB storage provided by the user?

A. 1 TB
B. 50 GB
C. 5 GB
D. 100 GB

Answer: D

Explanation:
If the user is trying to enable PIOPS with MySQL RDS, the minimum size of storage should be 100 GB. Reference:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html

NEW QUESTION 8
The Statement element, of an AWS IAM policy, contains an array of indMdual statements. Each indMdual statement is a(n) block enclosed in braces { }.

A. XML
B. JavaScript
C. JSON
D. AJAX

Answer: C

Explanation:
The Statement element, of an IAM policy, contains an array of indMdual statements. Each indMdual statement is a JSON block enclosed in braces { }.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPoIicyLanguage_EIementDescriptions.html

NEW QUESTION 9
If no explicit deny is found while applying IAM's Policy Evaluation Logic, the enforcement code looks for any instructions that would apply to the request.

A. "cancel"
B. "suspend"
C. "a||ow"
D. "vaIid"

Answer: C

Explanation:
If an explicit deny is not found among the applicable policies for a specific request, IAM's Policy Evaluation Logic checks for any "aIIow" instructions to check if the request can be successfully completed.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPoIicyLanguage_EvaIuationLogic.htmI

NEW QUESTION 10
A user has configured EBS volume with PIOPS. The user is not experiencing the optimal throughput. Which of the following could not be factor affecting I/O performance of that EBS volume?

A. EBS bandwidth of dedicated instance exceeding the PIOPS
B. EBS volume size
C. EC2 bandwidth
D. Instance type is not EBS optimized

Answer: B

Explanation:
If the user is not experiencing the expected IOPS or throughput that is provisioned, ensure that the EC2 bandwidth is not the limiting factor, the instance is EBS-optimized (or include 10 Gigabit network connectMty) and the instance type EBS dedicated bandwidth exceeds the IOPS more than he has provisioned.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-io-characteristics.html

NEW QUESTION 10
How can multiple compute resources be used on the same pipeline in AWS Data Pipeline?

A. You can use multiple compute resources on the same pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each actMty via its runsOn field.
B. You can use multiple compute resources on the same pipeline by defining multiple cluster definition files.
C. You can use multiple compute resources on the same pipeline by defining multiple clusters for your actMty.
D. You cannot use multiple compute resources on the same pipelin

Answer: A

Explanation:
Multiple compute resources can be used on the same pipeline in AWS Data Pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each actMty via its runsOn field, which allows pipelines to combine AWS and on-premise resources, or to use a mix of instance types for their actMties.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 15
The two policies that you attach to an IAM role are the access policy and the trust policy. The trust policy identifies who can assume the role and grants the permission in the AWS Lambda account principal by adding the action.

A. aws:AssumeAdmin
B. Iambda:InvokeAsync
C. sts:|nvokeAsync
D. sts:AssumeRoIe

Answer: D

Explanation:
The two policies that you attach to an IAM role are the access policy and the trust policy.
Remember that adding an account to the trust policy of a role is only half of establishing the trust relationship. By default, no users in the trusted accounts can assume the role until the administrator for that account grants the users the permission to assume the role by adding the Amazon Resource Name (ARN) of the role to an Allow element for the sts:AssumeRoIe action.
Reference: http://docs.aws.amazon.com/|AM/Iatest/UserGuide/id_ro|es_manage_modify.html

NEW QUESTION 20
You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

A. An XML document written in the IAM Policy Language
B. An XML document written in a language of your choice
C. A JSON document written in the IAM Policy Language
D. A JSON document written in a language of your choice

Answer: C

Explanation:
You define the permissions for a role in an IAM policy. An IAM policy is a JSON document written in the IAM Policy Language.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_ro|es_terms-and-concepts.html

NEW QUESTION 24
An organization is setting up an application on AWS to have both High Availabilty (HA) and Disaster Recovery (DR). The organization wants to have both Recovery point objective (RPO) and Recovery time objective (RTO) of 10 minutes. Which of the below mentioned service configurations does not help the organization achieve the said RPO and RTO?

A. Take a snapshot of the data every 10 minutes and copy it to the other region.
B. Use an elastic IP to assign to a running instance and use Route 53 to map the user's domain with that IP.
C. Create ELB with multi- region routing to allow automated failover when required.
D. Use an AMI copy to keep the AMI available in other region

Answer: C

Explanation:
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. Copy the AMI to another region to enable Disaster Recovery (DR) in case of region failure. The organization should also use EBS for persistent storage and take a snapshot every 10 minutes to meet Recovery time objective (RTO). They should also setup an elastic IP and use it with Route 53 to route requests to the same IP.
When one of the instances fails the organization can launch new instances and assign the same EIP to a new instance to achieve High Availability (HA). The ELB works only for a particular region and does not route requests across regions.
Reference: http://d36cz9buwru1tt.c|oudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 27
Does an AWS Direct Connect location provide access to Amazon Web Services in the region it is associated with as well as access to other US regions?

A. No, it provides access only to the region it is associated with.
B. No, it provides access only to the US regions other than the region it is associated with.
C. Yes, it provides access.
D. Yes, it provides access but only when there's just one Availability Zone in the regio

Answer: C

Explanation:
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. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCloud (US).
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

NEW QUESTION 28
Which of the following components of AWS Data Pipeline specifies the business logic of your data management?

A. Task Runner
B. Pipeline definition
C. AWS Direct Connect
D. Amazon Simple Storage Service (Amazon S3)

Answer: B

Explanation:
A pipeline definition specifies the business logic of your data management.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/what-is-datapipeline.htmI

NEW QUESTION 29
In IAM, which of the following is true of temporary security credentials?

A. Once you issue temporary security credentials, they cannot be revoked.
B. None of these are correct.
C. Once you issue temporary security credentials, they can be revoked only when the virtual MFA device is used.
D. Once you issue temporary security credentials, they can be revoke

Answer: A

Explanation:
Temporary credentials in IAM are valid throughout their defined duration of time and hence can't be revoked. However, because permissions are evaluated each time an AWS request is made using the credentials, you can achieve the effect of revoking the credentials by changing the permissions for the credentials even after they have been issued. Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentiaIs_temp_controI-access_disable-perms.h tml

NEW QUESTION 33
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": "AIIow", "Action": ["Describe"], "Resource": "BiIling"
B. "Effect": "AIIow", "Action": ["aws-portaI: ViewBi||ing"], "Resource": "*"
C. "Effect": "AIIow", "Action": ["aws-portaI:ViewUsage"], "Resource": "*"
D. "Effect": "AIIow", "Action": ["AccountUsage], "Resource": "*"

Answer: C

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": "A||ow", "Action": [
"aws-portaI:ViewUsage"
]!
"Resource": "*"
} I
}

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

NEW QUESTION 37
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

Explanation:
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 42
An organization is purchasing licensed software. The software license can be registered only to a specific MAC Address. The organization is going to host the software in the AWS environment. How can the organization fulfil the license requirement as the MAC address changes every time an instance is started/stopped/terminated?

A. It is not possible to have a fixed MAC address with AWS.
B. The organization should use VPC with the private subnet and configure the MAC address with that subnet
C. The organization should use VPC with an elastic network interface which will have a fixed MAC Address.
D. The organization should use VPC since VPC allows to configure the MAC address for each EC2 instance.

Answer: C

Explanation:
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. An Elastic Network Interface (ENI) is a virtual network interface that the user can attach to an instance in a VPC. An ENI can include attributes such as: a primary private IP address, one or more secondary private IP addresses, one elastic IP address per private IP address, one public IP address, one or more security groups, a MAC address, a source/destination check flag, and a description.
The user can create a network interface, attach it to an instance, detach it from an instance, and attach it to another instance. The attributes of a network interface follow the network interface as it is attached or detached from an instance and reattached to another instance. Thus, the user can maintain a fixed MAC using the network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.htmI

NEW QUESTION 43
An organization is undergoing a security audit. The auditor wants to view the AWS VPC configurations as the organization has hosted all the applications in the AWS VPC. The auditor is from a remote place and wants to have access to AWS to view all the VPC records.
How can the organization meet the expectations of the auditor without compromising on the security of their AWS infrastructure?

A. The organization should not accept the request as sharing the credentials means compromising on security.
B. Create an IAM role which will have read only access to all EC2 services including VPC and assign that role to the auditor.
C. Create an IAM user who will have read only access to the AWS VPC and share those credentials with the auditor.
D. The organization should create an IAM user with VPC full access but set a condition that will not allow to modify anything if the request is from any IP other than the organization's data center.

Answer: C

Explanation:
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 works with IAM and the organization can create IAM users who have access to various VPC services.
If an auditor wants to have access to the AWS VPC to verify the rules, the organization should be careful before sharing any data which can allow making updates to the AWS infrastructure. In this scenario it is recommended that the organization creates an IAM user who will have read only access to the VPC. Share the above mentioned credentials with the auditor as it cannot harm the organization. The sample policy is given below:
{
"Effect":"AI|ow",
"Action":["ec2:DescribeVpcs", "ec2:DescribeSubnets",
"ec2:DescribeInternetGateways", "ec2:DescribeCustomerGateways", "ec2:DescribeVpnGateways", "ec2:DescribeVpnConnections", "ec2:DescribeRouteTabIes",
"ec2:DescribeAddresses", "ec2:DescribeSecurityGroups", "ec2:DescribeNetworkAcIs", "ec2:DescribeDhcpOptions", "ec2:DescribeTags", "ec2:DescribeInstances"
]!
"Resource":"*"
}
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_IANI.htmI

NEW QUESTION 48
In Amazon Cognito, your mobile app authenticates with the Identity Provider (|dP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new for the user and a set of temporary, limited-prMlege AWS credentials.

A. Cognito Key Pair
B. Cognito API
C. Cognito ID
D. Cognito SDK

Answer: C

Explanation:
Your mobile app authenticates with the identity provider (IdP) using the provider's SDK. Once the end user is authenticated with the IdP, the OAuth or OpenID Connect token returned from the IdP is passed by your app to Amazon Cognito, which returns a new Cognito ID for the user and a set of temporary, limited-prMlege AWS credentials.
Reference: http://aws.amazon.com/cognito/faqs/

NEW QUESTION 51
What is the maximum length for a certificate ID in AWS IAM?

A. 1024 characters
B. 512 characters
C. 64 characters
D. 128 characters

Answer: D

Explanation:
The maximum length for a certificate ID is 128 characters.
Reference: http://docs.aws.amazon.com/IANI/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 52
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 indMdual 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 user
E. All the users should use their own keys to connect with the instance securely.

Answer: A

Explanation:
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 prMleges, such as sudo for users who need to perform occasional prMleged 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.
Reference: http://aws.amazon.com/articles/1233/

NEW QUESTION 55
An organization is setting up a highly scalable application using Elastic Beanstalk. They are using Elastic Load Balancing (ELB) as well as a Virtual Private Cloud (VPC) with public and private subnets. They have the following requirements:
. All the EC2 instances should have a private IP
. All the EC2 instances should receive data via the ELB's. Which of these will not be needed in this setup?

A. Launch the EC2 instances with only the public subnet.
B. Create routing rules which will route all inbound traffic from ELB to the EC2 instances.
C. Configure ELB and NAT as a part of the public subnet only.
D. Create routing rules which will route all outbound traffic from the EC2 instances through NA

Answer: A

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. If the organization wants the Amazon EC2 instances to have a private IP address, he should create a public and private subnet for VPC in each Availability Zone (this is an AWS Elastic Beanstalk requirement). The organization should add their public resources, such as ELB and NAT to the public subnet, and AWS Elastic Beanstalk will assign them unique elastic IP addresses (a static, public IP address). The organization should launch Amazon EC2 instances in a private subnet so that AWS Elastic Beanstalk assigns them non-routable private IP addresses. Now the organization should configure route tables with the following rules:
. route all inbound traffic from ELB to EC2 instances
. route all outbound traffic from EC2 instances through NAT
Reference: http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/AWSHowTo-vpc.html

NEW QUESTION 59
An EC2 instance that performs source/destination checks by default is launched in a private VPC subnet. All security, NACL, and routing definitions are configured as expected. A custom NAT instance is launched.
Which of the following must be done for the custom NAT instance to work?

A. The source/destination checks should be disabled on the NAT instance.
B. The NAT instance should be launched in public subnet.
C. The NAT instance should be configured with a public IP address.

D. The NAT instance should be configured with an elastic IP addres

Answer: A

Explanation:
Each EC2 instance performs source/destination checks by default. This means that the instance must be the source or destination of any traffic it sends or receives. However, a NAT instance must be able to send and receive traffic when the source or destination is not itself. Therefore, you must disable source/destination checks on the NAT instance.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/UserGuide/VPC_NAT_|nstance.htm|#EIP_Disab|e_Src DestCheck

NEW QUESTION 62
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. StringEquaIsIgnoreCase
B. StringNotEquaIsIgnoreCase
C. StringLikeStringEqua|s
D. StringNotEqua|s

Answer: A

Explanation:
When using string conditions within IANI, short versions of the available comparators can be used instead of the more verbose versions. For instance, streqi is the short version of StringEqua|s|gnoreCase that checks for the exact match between two strings ignoring their case.
Reference: http://awsdocs.s3.amazonaws.com/SNS/20100331/sns-gsg-2010-03-31.pdf

NEW QUESTION 65
Attempts, one of the three types of items associated with the schedule pipeline in the AWS Data Pipeline, provides robust data management.
Which of the following statements is NOT true about Attempts?

A. Attempts provide robust data management.
B. AWS Data Pipeline retries a failed operation until the count of retries reaches the maximum number of allowed retry attempts.
C. An AWS Data Pipeline Attempt object compiles the pipeline components to create a set of actionable instances.
D. AWS Data Pipeline Attempt objects track the various attempts, results, and failure reasons if applicable.

Answer: C

Explanation:
Attempts, one of the three types of items associated with a schedule pipeline in AWS Data Pipeline, provides robust data management. AWS Data Pipeline retries a failed operation. It continues to do so until the task reaches the maximum number of allowed retry attempts. Attempt objects track the various attempts, results, and failure reasons if applicable. Essentially, it is the instance with a counter. AWS Data Pipeline performs retries using the same resources from the previous attempts, such as Amazon EMR clusters and EC2 instances.
Reference:
http://docs.aws.amazon.com/datapipeline/latest/DeveIoperGuide/dp-how-tasks-scheduled.htmI

NEW QUESTION 69
In Amazon RDS for PostgreSQL, you can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xIarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve:

A. higher latency and lower throughput.
B. lower latency and higher throughput.
C. higher throughput only.
D. higher latency onl

Answer: B

Explanation:
You can provision up to 3TB storage and 30,000 IOPS per database instance. For a workload with 50% writes and 50% reads running on a cr1.8xIarge instance, you can realize over 25,000 IOPS for PostgreSQL. However, by provisioning more than this limit, you may be able to achieve lower latency and higher throughput. Your actual realized IOPS may vary from the amount you provisioned based on your database workload, instance type, and database engine choice.
Reference: https://aws.amazon.com/rds/postgresq|/

NEW QUESTION 71
Within an IAM policy, can you add an IfExists condition at the end of a Null condition?

A. Yes, you can add an IfExists condition at the end of a Null condition but not in all Regions.
B. Yes, you can add an IfExists condition at the end of a Null condition depending on the condition.
C. No, you cannot add an IfExists condition at the end of a Null condition.
D. Yes, you can add an IfExists condition at the end of a Null conditio

Answer: C

Explanation:
Within an IAM policy, IfExists can be added to the end of any condition operator except the Null condition. It can be used to indicate that conditional comparison needs to happen if the policy key is present in the context of a request; otherwise, it can be ignored.
Reference: http://docs.aws.amazon.com/IAM/Iatest/UserGuide/reference_poIicies_eIements.html

NEW QUESTION 73
IAM users do not have permission to create Temporary Security Credentials for federated users and roles by default. In contrast, IAM users can call without the need of any special permissions

A. GetSessionName
B. GetFederationToken
C. GetSessionToken
D. GetFederationName

Answer: C

Explanation:
Currently the STS API command GetSessionToken is available to every IAM user in your account without previous permission. In contrast, the GetFederationToken command is restricted and explicit permissions need to be granted so a user can issue calls to this particular Action
Reference: http://docs.aws.amazon.com/STS/latest/UsingSTS/STSPermission.html

NEW QUESTION 74
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 subne

Answer: A

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. Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/UserGuide/VPC_Introduction.htm|

NEW QUESTION 75
An organization has developed an application which provides a smarter shopping experience. They need to show a demonstration to various stakeholders who may not be able to access the in premise
application so they decide to host a demo version of the application on AWS. Consequently they will need a fixed elastic IP attached automatically to the instance when it is launched.
In this scenario which of the below mentioned options will not help assign the elastic IP automatically?

A. Write a script which will fetch the instance metadata on system boot and assign the public IP using that metadata.
B. Provide an elastic IP in the user data and setup a bootstrapping script which will fetch that elastic IP and assign it to the instance.
C. Create a controlling application which launches the instance and assigns the elastic IP based on the parameter provided when that instance is booted.
D. Launch instance with VPC and assign an elastic IP to the primary network interfac

Answer: A

Explanation:
EC2 allows the user to launch On-Demand instances. If the organization is using an application temporarily only for demo purposes the best way to assign an elastic IP would be:
Launch an instance with a VPC and assign an EIP to the primary network interface. This way on every instance start it will have the same IP Create a bootstrapping script and provide it some metadata, such as user data which can be used to assign an EIP Create a controller instance which can schedule the start and stop of the instance and provide an EIP as a parameter so that the controller instance can check the instance boot and assign an EIP
The instance metadata gives the current instance data, such as the public/private IP. It can be of no use for assigning an EIP.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-instancedata.html

NEW QUESTION 76
An organization is having a VPC for the HR department, and another VPC for the Admin department. The HR department requires access to all the instances running in the Admin VPC while the Admin department requires access to all the resources in the HR department. How can the organization setup this scenario?

A. Setup VPC peering between the VPCs of Admin and HR.
B. Setup ACL with both VPCs which will allow traffic from the CIDR of the other VPC.
C. Setup the security group with each VPC which allows traffic from the CIDR of another VPC.
D. It is not possible to connect resources of one VPC from another VPC.

Answer: A

Explanation:
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.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html

NEW QUESTION 77
ExamKiller 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. ExamKiller wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ExamKiller also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons. How can ExamKil|er setup this

scenario?

A. ExamKiI|er has to setup one centralized VPC which will peer in to all the other VPCs of the tenants.
B. ExamKiI|er 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. ExamKiI|er should setup all the VPCs with the same CIDR but have a centralized VP
D. This way only the centralized VPC can talk to the other VPCs using VPC peering.
E. ExamKiI|er should setup all the VPCs meshed together with VPC peering for all VPC

Answer: A

Explanation:
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.
Reference:
http://docs.aws.amazon.com/AmazonVPC/Iatest/PeeringGuide/peering-configurations-full-access.html# many-vpcs-full-acces

NEW QUESTION 79
In Amazon Redshift, how many slices does a dw2.8xIarge node have?

A. 16
B. 8
C. 32
D. 2

Answer: C

Explanation:
The disk storage for a compute node in Amazon Redshift is dMded into a number of slices, equal to the number of processor cores on the node. For example, each DW1.XL compute node has two slices, and each DW2.8XL compute node has 32 slices.
Reference: http://docs.aws.amazon.com/redshift/latest/dg/t_Distributing_data.html

NEW QUESTION 80
Identify a true statement about using an IAM role to grant permissions to applications running on Amazon EC2 instances.

A. When AWS credentials are rotated, developers have to update only the root Amazon EC2 instance that uses their credentials.
B. When AWS credentials are rotated, developers have to update only the Amazon EC2 instance on which the password policy was applied and which uses their credentials.
C. When AWS credentials are rotated, you don't have to manage credentials and you don't have to worry about long-term security risks.
D. When AWS credentials are rotated, you must manage credentials and you should consider precautions for long-term security risks.

Answer: C

Explanation:
Using IAM roles to grant permissions to applications that run on EC2 instances requires a bit of extra configuration. Because role credentials are temporary and rotated automatically, you don't have to manage credentials, and you don't have to worry about long-term security risks.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.html

NEW QUESTION 82
Out of the striping options available for the EBS volumes, which one has the following disadvantage: 'Doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.'?

A. Raid 1
B. Raid 0
C. RAID 1+0 (RAID 10)
D. Raid 2

Answer: C

Explanation:
RAID 1+0 (RAID 10) doubles the amount of I/O required from the instance to EBS compared to RAID 0, because you're mirroring all writes to a pair of volumes, limiting how much you can stripe.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

NEW QUESTION 84
In the context of IAM roles for Amazon EC2, which of the following NOT true about delegating permission to make API requests?

A. You cannot create an IAM role.
B. You can have the application retrieve a set of temporary credentials and use them.
C. You can specify the role when you launch your instances.
D. You can define which accounts or AWS services can assume the rol

Answer: A

Explanation:

Amazon designed IANI roles so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles as follows: Create an IAM role. Define which accounts or AWS services can assume the role. Define which API actions and resources the application can use after assuming the role. Specify the role when you launch your instances. Have the application retrieve a set of temporary credentials and use them.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

NEW QUESTION 85
In the context of Amazon E|astiCache CLI, which of the following commands can you use to view all EIastiCache 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

Explanation:
In Amazon EIastiCache, 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). Reference:
http://docs.aws.amazon.com/AmazonEIastiCache/Iatest/UserGuide/ECEvents.Viewing.html

NEW QUESTION 90
In Amazon Cognito what is a silent push notification?

A. It is a push message that is received by your application on a user's device that will not be seen by theusen
B. It is a push message that is received by your application on a user's device that will return the user's geolocation.
C. It is a push message that is received by your application on a user's device that will not be heard by the usen
D. It is a push message that is received by your application on a user's device that will return the user's authentication credentials.

Answer: A

Explanation:
Amazon Cognito uses the Amazon Simple Notification Service (SNS) to send silent push notifications to devices. A silent push notification is a push message that is received by your application on a user's device that will not be seen by the user.
Reference: http://aws.amazon.com/cognito/faqs/

NEW QUESTION 91
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 "NumericLessThanEquaIs"?

A. numlteq
B. numlteql
C. numltequals
D. numeql

Answer: A

Explanation:
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 NumericLessThanEquaIs.
Reference: http://awsdocs.s3.amazonaws.com/SQS/2011-10-01/sqs-dg-2011-10-01.pdf

NEW QUESTION 93
AWS has launched T2 instances which come with CPU usage credit. An organization has a requirement which keeps an instance running for 24 hours. However, the organization has high usage only during 11 AM to 12 PM. The organization is planning to use a T2 small instance for this purpose.
If the organization already has multiple instances running since Jan 2012, which of the below mentioned options should the organization implement while launching a T2 instance?

A. The organization must migrate to the EC2-VPC platform first before launching a T2 instance.
B. While launching a T2 instance the organization must create a new AWS account as this account does not have the EC2-VPC platform.
C. Create a VPC and launch a T2 instance as part of one of the subnets of that VPC.
D. While launching a T2 instance the organization must select EC2-VPC as the platform.

Answer: C

Explanation:
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 AWS account provides two platforms:
EC2-CLASSIC and EC2-VPC, depending on when the user has created his AWS account and which regions he is using. If the user has created the AWS account after 2013-12-04, it supports only EC2-VPC. In this scenario, since the account is before the required date the supported platform will be
EC2-CLASSIC. It is required that the organization creates a VPC as the T2 instances can be launched only as a part of VPC.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/vpc-migrate.html

NEW QUESTION 94
How does AWS Data Pipeline execute actMties on on-premise resources or AWS resources that you manage?

A. By supplying a Task Runner package that can be installed on your on-premise hosts
B. None of these

C. By supplying a Task Runner file that the resources can access for execution
D. By supplying a Task Runnerjson script that can be installed on your on-premise hosts

Answer: A

Explanation:
To enable running actMties using on-premise resources, AWS Data Pipeline does the following: It supply a Task Runner package that can be installed on your on-premise hosts.
This package continuously polls the AWS Data Pipeline service for work to perform.
When it's time to run a particular actMty on your on-premise resources, it will issue the appropriate command to the Task Runner.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 99
A user is configuring MySQL RDS with PIOPS. What should be the minimum PIOPS that the user should provision?

A. 1000
B. 200
C. 2000
D. 500

Answer: A

Explanation:
If a user is trying to enable PIOPS with MySQL RDS, the minimum size of storage should be 100 GB and the minimum PIOPS should be 1000.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.html

NEW QUESTION 102
Do you need to use Amazon Cognito to use the Amazon Mobile Analytics service?

A. N
B. However, it is recommend by AWS to use Amazon Cognito for security best practices.
C. Ye
D. You need to use it only if you have IAM root access.
E. N
F. You cannot use it at all, and you need to use AWS IAM accounts.
G. Ye
H. It is recommended by AWS to use Amazon Cognito to use Amazon Mobile Analytics servic

Answer: A

Explanation:
You can initialize Amazon Mobile Analytics using AWS IAM accounts. AWS recommend using Amazon Cognito for security best practices.
Reference: http://aws.amazon.com/mobi|eanaIytics/faqs/

NEW QUESTION 105
You want to use Amazon Redshift and you are planning to deploy dw1.8xIarge nodes. What is the minimum amount of nodes that you need to deploy with this kind of configuration?

A. 1
B. 4
C. 3
D. 2

Answer: D

Explanation:
For a single-node configuration in Amazon Redshift, the only option available is the smallest of the two options. The 8XL extra-large nodes are only available in a multi-node configuration
Reference: http://docs.aws.amazon.com/redshift/latest/mgmt/working-with-c|usters.htmI

NEW QUESTION 107
To get started using AWS Direct Connect, in which of the following steps do you configure Border Gateway Protocol (BGP)?

A. Complete the Cross Connect
B. Configure Redundant Connections with AWS Direct Connect
C. Create a Virtual Interface
D. Download Router Configuration

Answer: C

Explanation:
In AWS Direct Connect, your network must support Border Gateway Protocol (BGP) and BGP MD5 authentication, and you need to provide a private Autonomous System Number (ASN) for that to connect to Amazon Virtual Private Cloud (VPC). To connect to public AWS products such as Amazon EC2 and Amazon S3, you will also need to provide a public ASN that you own (preferred) or a private ASN. You have to configure BGP in the Create a Virtual Interface step.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/getstarted.htmI#createvirtualinterface

NEW QUESTION 108

Which of the following components of AWS Data Pipeline polls for tasks and then performs those tasks?

A. Pipeline Definition
B. Task Runner
C. Amazon Elastic MapReduce (EMR)
D. AWS Direct Connect

Answer: B

Explanation:
Task Runner polls for tasks and then performs those tasks.
Reference: http://docs.aws.amazon.com/datapipeline/latest/DeveloperGuide/what-is-datapipeline.html

NEW QUESTION 111
A user is hosting a public website on AWS. The user wants to have the database and the app server on the AWS VPC. The user wants to setup a database that can connect to the Internet for any patch upgrade but cannot receive any request from the internet. How can the user set this up?

A. Setup DB in a private subnet with the security group allowing only outbound traffic.
B. Setup DB in a public subnet with the security group allowing only inbound data.
C. Setup DB in a local data center and use a private gateway to connect the application with DB.
D. Setup DB in a private subnet which is connected to the internet via NAT for outbound.

Answer: D

Explanation:
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. AWS provides two features that the user can use to increase security in VPC: security groups and network ACLs. When the user wants to setup both the DB and App on VPC, the user should make one public and one private subnet. The DB should be hosted in a private subnet and instances in that subnet cannot reach the internet. The user can allow an instance in his VPC to initiate outbound connections to the internet but prevent unsolicited inbound connections from the internet by using a Network Address Translation (NAT) instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION 114
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

Explanation:
CloudWatch is a monitoring tool and doesn't give users access to manage Amazon ElastiCache. Reference:
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.NIanaging.html

NEW QUESTION 117
An organization is hosting a scalable web application using AWS. The organization has configured internet facing ELB and Auto Scaling to make the application scalable. Which of the below mentioned
statements is required to be followed when the application is planning to host a web application on VPC?

A. The ELB can be in a public or a private subnet but should have the ENI which is attached to an elastic IP.
B. The ELB must not be in any subnet; instead it should face the internet directly.
C. The ELB must be in a public subnet of the VPC to face the internet traffic.
D. The ELB can be in a public or a private subnet but must have routing tables attached to divert the internet traffic to it.

Answer: C

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 ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For internet facing ELB it is required that ELB should be in a public subnet.
After the user creates the public subnet, he should ensure to associate the route table of the public subnet with the internet gateway to enable the load balancer in the subnet to connect with the internet. Reference: http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/CreateVPCForELB.html

NEW QUESTION 122
Your company has recently extended its datacenter into a VPC on AVVS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

A. Use OAuth 2 0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS IV|anagement Console.
C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
D. Use your on-premises SAML 2.0-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: D

NEW QUESTION 125
You're running an application on-premises due to its dependency on non-x86 hardware and want to use AWS for data backup. Your backup application is only able to write to POSIX-compatible block-based storage. You have 140TB of data and would like to mount it as a single folder on your file server Users must be able to access portions of this data while the backups are taking place. What backup solution would be most appropriate for this use case?

A. Use Storage Gateway and configure it to use Gateway Cached volumes.
B. Configure your backup software to use S3 as the target for your data backups.
C. Configure your backup software to use Glacier as the target for your data backups.
D. Use Storage Gateway and configure it to use Gateway Stored volume

Answer: A

NEW QUESTION 127
You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. Dunning a DR test you notice that when you disable all web sewers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

A. Latency resource record sets cannot be used in combination with weighted resource record sets.
B. You did not setup an HTTP health check to one or more of the weighted resource record sets associated with me disabled web sewers.
C. The value of the weight associated with the latency alias resource record set in the region with the disabled sewers is higher than the weight for the other region.
D. One of the two working web sewers in the other region did not pass its HTTP health check.
E. You did not set "Evaluate Target Health" to "Yes" on the latency alias resource record set associated with example com in the region where you disabled the servers.

Answer: BE

NEW QUESTION 132
Your startup wants to implement an order fulfillment process for selling a personalized gadget that needs an average of 3-4 days to produce with some orders taking up to 6 months you expect 10 orders per day on your first day. 1000 orders per day after 6 months and 10,000 orders after 12 months.
Orders coming in are checked for consistency men dispatched to your manufacturing plant for production quality control packaging shipment and payment processing If the product does not meet the quality standards at any stage of the process employees may force the process to repeat a step Customers are notified via email about order status and any critical issues with their orders such as payment failure.
Your case architecture includes AWS Elastic Beanstalk for your website with an RDS MySQL instance for customer data and orders.
How can you implement the order fulfillment process while making sure that the emails are delivered reliably?

A. Add a business process management application to your Elastic Beanstalk app sewers and re-use the ROS database for tracking order status use one of the Elastic Beanstalk instances to send emails to customers.
B. Use SWF with an Auto Scaling group of actMty workers and a decider instance in another Auto Scaling group with min/max=1 Use the decider instance to send emails to customers.
C. Use SWF with an Auto Scaling group of actMty workers and a decider instance in another Auto Scaling group with min/max=1 use SES to send emails to customers.
D. Use an SQS queue to manage all process tasks Use an Auto Scaling group of EC2 Instances that poll the tasks and execute the
E. Use SES to send emails to customers.

Answer: C

NEW QUESTION 135
You have a periodic Image analysis application that gets some files In Input analyzes them and tor each file writes some data in output to a ten file the number of files in input per day is high and concentrated in a few hours of the day.
Currently you have a server on EC2 with a large EBS volume that hosts the input data and the results it takes almost 20 hours per day to complete the process
What services could be used to reduce the elaboration time and improve the availability of the solution?

A. S3 to store I/O file
B. SQS to distribute elaboration commands to a group of hosts working in paralle
C. Auto scaling to dynamically size the group of hosts depending on the length of the SQS queue
D. EBS with Provisioned IOPS (PIOPS) to store I/O file
E. SNS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group of hosts depending on the number of SNS notifications
F. S3 to store I/O files, SNS to distribute evaporation commands to a group of hosts working in paralle
G. Auto scaling to dynamically size the group of hosts depending on the number of SNS notifications
H. EBS with Provisioned IOPS (PIOPS) to store I/O files SQS to distribute elaboration commands to a group of hosts working in parallel Auto Scaling to dynamically size the group ot hosts depending on the length of the SQS queue.

Answer: D

NEW QUESTION 139
You have been asked to design the storage layer for an application. The application requires disk performance of at least 100,000 IOPS. In addition, the storage layer must be able to survive the loss of an indMdual disk, EC2 instance, or Availability Zone without any data loss. The volume you provide must have a capacity of at least 3 TB. Which of the following designs will meet these objectives?

A. Instantiate a c3.8x|arge instance in us-east-1. Provision 4x1TB EBS volumes, attach them to the instance, and configure them as a single RAID 5 volum
B. Ensure that EBS snapshots are performed every 15 minutes.
C. Instantiate a c3.8xIarge instance in us-east-1. Provision 3xITB EBS volumes, attach them to the Instance, and configure them as a single RAID 0 volum
D. Ensure that EBS snapshots are performed every 15 minutes.
E. Instantiate an i2.8xIarge instance in us-east-1
F. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instanc

G. Provision 3x1TB EBS volumes, attach them to the instance, and configure them as a second RAID 0 volum
H. Configure synchronous, block-level replication from the ephemeral-backed volume to the EBS-backed volume.
I. Instantiate a c3.8xIarge instance in us-east-1. Provision an AWS Storage Gateway and configure it for 3 TB of storage and 100,000 IOP
J. Attach the volume to the instance.
K. Instantiate an i2.8xIarge instance in us-east-1
L. Create a RAID 0 volume using the four 800GB SSD ephemeral disks provided with the instanc
M. Configure synchronous, blocklevel replication to an identically configured instance in us-east-1b.

Answer: C

NEW QUESTION 144
You are the new IT architect in a company that operates a mobile sleep tracking application.
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.
The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app.
Currently you have around 100k users who are mostly based out of North America. You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? Choose 2 answers

A. Have the mobile app access Amazon DynamoDB directly Instead of JSON files stored on Amazon S3.
B. Write data directly into an Amazon Redshift cluster replacing both Amazon DynamoDB and Amazon S3.
C. Introduce an Amazon SQS queue to buffer writes to the Amazon DynamoDB table and reduce provisioned write throughput.
D. Introduce Amazon Elasticache to cache reads from the Amazon DynamoDB table and reduce provisioned read throughput.
E. Create a new Amazon DynamoDB table each day and drop the one for the previous day after its data is on Amazon S3.

Answer: AD

NEW QUESTION 145
A large real-estate brokerage is exploring the option o(adding a cost-effective location based alert to their existing mobile application The application backend infrastructure currently runs on AWS Users who opt in to this service will receive alerts on their mobile device regarding real-estate otters in proximity to their location. For the alerts to be relevant delivery time needs to be in the low minute count the existing mobile app has 5 million users across the US. Which one of the following architectural suggestions would you make to the customer?

A. The mobile application will submit its location to a web service endpoint utilizing Elastic Load Balancing and EC2 instances: DynamoDB will be used to store and retrieve relevant offers EC2 instances will communicate with mobile earners/device providers to push alerts back to mobile application.
B. Use AWS DirectConnect or VPN to establish connectMty with mobile carriers EC2 instances will receive the mobile applications ' location through carrier connection: RDS will be used to store and relevant offers EC2 instances will communicate with mobile carriers to push alerts back to the mobile application
C. The mobile application will send device location using SQ
D. EC2 instances will retrieve the relevant others from DynamoDB AWS MobiIe Push will be used to send offers to the mobile application
E. The mobile application will send device location using AWS Nlobile Push EC2 instances will retrieve the relevant offers from DynamoDB EC2 instances will communicate with mobile carriers/device providers to push alerts back to the mobile application.

Answer: A

NEW QUESTION 150
Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of the various servers. The backup system must support database recovery whole server and whole disk restores, and indMdual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database
Which backup architecture will meet these requirements?

A. Backup RDS using automated daily DB backups Backup the EC2 instances using AMIs and supplement with file-level backup to S3 using traditional enterprise backup software to provide file level restore
B. Backup RDS using a Multi-AZ Deployment Backup the EC2 instances using Amis, and supplement by copying file system data to S3 to provide file level restore.
C. Backup RDS using automated daily DB backups Backup the EC2 instances using EBS snapshots and supplement with file-level backups to Amazon Glacier using traditional enterprise backup software to provide file level restore
D. Backup RDS database to S3 using Oracle RMAN Backup the EC2 instances using Amis, and supplement with EBS snapshots for indMdual volume restore.

Answer: A

NEW QUESTION 154
An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least prMlege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party.
Which of the following would meet all of these conditions?

A. From the AWS Management Console, navigate to the Security Credentials page and retrieve the access and secret key for your account.
B. Create an IAM user within the enterprise account assign a user policy to the IAM user that allows only the actions required by the SaaS application create a new access and secret key for the user and provide these credentials to the SaaS provider.
C. Create an IAM role for cross-account access allows the SaaS provider's account to assume the role and assign it a policy that allows only the actions required by the SaaS application.
D. Create an IAM role for EC2 instances, assign it a policy that allows only the actions required tor the Saas application to work, provide the role ARN to the SaaS provider to use when launching their application instances.

Answer: C

NEW QUESTION 157
An AWS customer runs a public blogging website. The site users upload two million blog entries a month. The average blog entry size is 200 KB. The access rate to blog entries drops to negligible 6 months after publication and users rarely access a blog entry 1 year after publication. Additionally, blog entries have a high

update rate during the first 3 months following publication, this drops to no updates after 6 months. The customer wants to use CloudFront to improve his user's load times. Which of the following recommendations would you make to the customer?

A. Duplicate entries into two different buckets and create two separate CloudFront distributions where S3 access is restricted only to Cloud Front identity
B. Create a CloudFront distribution with "US Europe" price class for US/Europe users and a different CloudFront distribution with "AII Edge Locations" for the remaining users.
C. Create a CloudFront distribution with S3 access restricted only to the CloudFront identity and partition the blog entry's location in S3 according to the month it was uploaded to be used with CloudFront behaviors.
D. Create a CloudFront distribution with Restrict Viewer Access Forward Query string set to true and minimum TTL of 0.

Answer: C

NEW QUESTION 161
You are implementing a URL whitelisting system for a company that wants to restrict outbound HTTP'S connections to specific domains from their EC2-hosted applications you deploy a single EC2 instance running proxy software and configure It to accept traffic from all subnets and EC2 instances in the VPC. You configure the proxy to only pass through traffic to domains that you define in its whitelist configuration You have a nightly maintenance window or 10 minutes where ail instances fetch new software updates. Each update Is about 200MB In size and there are 500 instances In the VPC that routinely fetch updates After a few days you notice that some machines are failing to successfully download some, but not all of their updates within the maintenance window. The download URLs used for these updates are correctly listed in the proxy's whitelist configuration and you are able to access them manually using a web browser on the instances. What might be happening? (Choose 2 answers)

A. You are running the proxy on an undersized EC2 instance type so network throughput is not sufficient for all instances to download their updates in time.
B. You are running the proxy on a sufficiently-sized EC2 instance in a private subnet andits network throughput is being throttled by a NAT running on an undersized EC2 instance.
C. The route table for the subnets containing the affected EC2 instances is not configured to direct network traffic for the software update locations to the proxy.
D. You have not allocated enough storage to the EC2 instance running the proxy so the network buffer is filling up, causing some requests to fail.
E. You are running the proxy in a public subnet but have not allocated enough EIPs to support the needed network throughput through the Internet Gateway (IGW).

Answer: AB

NEW QUESTION 162
You are designing a social media site and are considering how to mitigate distributed denial-of-service (DDoS) attacks. Which of the below are viable mitigation techniques? (Choose 3 answers)

A. Add multiple elastic network interfaces (ENIs) to each EC2 instance to increase the network bandwidth.
B. Use dedicated instances to ensure that each instance has the maximum performance possible.
C. Use an Amazon C|oudFront distribution for both static and dynamic content.
D. Use an Elastic Load Balancer with auto scaling groups at the we
E. App and Amazon Relational Database Service (RDS) tiers
F. Add alert Amazon CloudWatch to look for high Network in and CPU utilization.
G. Create processes and capabilities to quickly add and remove rules to the instance OS firewal

Answer: CEF

NEW QUESTION 164
Your website is serving on-demand training videos to your workforce. Videos are uploaded monthly in high resolution MP4 format. Your workforce is distributed globally often on the move and using company-provided tablets that require the HTTP Live Streaming (HLS) protocol to watch a video. Your company has no video transcoding expertise and it required you may need to pay for a consultant.
How do you implement the most cost-efficient architecture without compromising high availability and
quality of video delivery'?

A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queu
B. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
C. CloudFront to serve HLS transcoded videos from EC2.
D. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
E. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few day
F. CloudFront to serve HLS transcoded videos from EC2.
G. Elastic Transcoder to transcode original high-resolution MP4 videos to HL
H. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few day
I. C|oudFront to serve HLS transcoded videos from S3.
J. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queu
K. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few day
L. CloudFront to serve HLS transcoded videos from Glacier.

Answer: C

NEW QUESTION 168
A customer has established an AWS Direct Connect connection to AWS. The link is up and routes are being advertised from the customer's end, however the customer is unable to connect from EC2 instances inside its VPC to servers residing in its datacenter.
Which of the following options provide a viable solution to remedy this situation? (Choose 2 answers)

A. Add a route to the route table with an IPsec VPN connection as the target.
B. Enable route propagation to the virtual pinnate gateway (VGW).
C. Enable route propagation to the customer gateway (CGW).
D. Modify the route table of all Instances using the 'route' command.
E. Modify the Instances VPC subnet route table by adding a route back to the customer's on-premises environment.

Answer: AC

NEW QUESTION 173
You are running a news website in the eu-west-1 region that updates every 15 minutes. The website has a world-wide audience it uses an Auto Scaling group behind an Elastic Load Balancer and an Amazon
RDS database Static content resides on Amazon S3, and is distributed through Amazon CIoudFront. Your Auto Scaling group is set to trigger a scale up event at 60% CPU utilization, you use an Amazon RDSextra large DB instance with 10.000 Provisioned IOPS its CPU utilization is around 80%. While freeable memory is in the 2 GB range.
Web analytics reports show that the average load time of your web pages is around 1.5 to 2 seconds, but your SEO consultant wants to bring down the average load time to under 0.5 seconds.
How would you improve page load times for your users? (Choose 3 answers)

A. Lower the scale up trigger of your Auto Scaling group to 30% so it scales more aggressively.
B. Add an Amazon EIastiCache caching layer to your application for storing sessions and frequent DB quenes
C. Configure Amazon CIoudFront dynamic content support to enable caching of re-usable content from your site
D. Switch the Amazon RDS database to the high memory extra large Instance type
E. Set up a second installation in another region, and use the Amazon Route 53 latency-based routing feature to select the right region.

Answer: ABD

NEW QUESTION 178
Your company hosts a social media website for storing and sharing documents. The web application allows user to upload large files while resuming and pausing the upload as needed. Currently, files are uploaded to your PHP front end backed by Elastic load Balancing and an autoscaling fileet of Amazon Elastic Compute Cloud (EC2) instances that scale upon average of bytes received (Networkln). After a file has been uploaded, it is copied to Amazon Simple Storage Service (S3). Amazon EC2 instances use an AWS Identity and Access Management (IAM) role that allows Amazon S3 uploads. Over the last six months, your user base and scale have increased significantly, forcing you to increase the Auto Scaling group's Max parameter a few times. Your CFO is concerned about rising costs and has asked you to adjust the architecture where needed to better optimize costs.
Which architecture change could you introduce to reduce costs and still keep your web application secure and scalable?

A. Replace the Auto Scaling launch configuration to include c3.8xIarge instances; those instances can potentially yield a network throuthput of 10gbps.
B. Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your ap
C. Implement client-side logic to directly upload the file to Amazon S3 using the given credentials and S3 prefix.
D. Re-architect your ingest pattern, and move your web application instances into a VPC public subne
E. Attach a public IP address for each EC2 instance (using the Auto Scaling launch configuration settings). Use Amazon Route 53 Round Robin records set and HTTP health check to DNS load balance the apprequests; this approach will significantly reduce the cost by bypassing Elastic Load Balancing.
F. Re-architect your ingest pattern, have the app authenticate against your identity provider, and use your identity provider as a broker fetching temporary AWS credentials from AWS Secure Token Service (GetFederationToken). Securely pass the credentials and S3 endpoint/prefix to your ap
G. Implement client-side logic that used the S3 multipart upload API to directly upload the file to Amazon S3 using the given credentials and S3 prefix.

Answer: C

NEW QUESTION 180
You require the ability to analyze a customer's clickstream data on a website so they can do behavioral analysis. Your customer needs to know what sequence of pages and ads their customer clicked on. This data will be used in real time to modify the page layouts as customers click through the site to increase stickiness and advertising click-through. Which option meets the requirements for captioning and analyzing this data?

A. Log clicks in weblogs by URL store to Amazon S3, and then analyze with Elastic MapReduce
B. Push web clicks by session to Amazon Kinesis and analyze behavior using Kinesis workers
C. Write click events directly to Amazon Redshift and then analyze with SQL
D. Publish web clicks by session to an Amazon SQS queue then periodically drain these events to Amazon RDS and analyze with SQL.

Answer: B

NEW QUESTION 184
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 sewers, two database sewers and one NAT instance tor a total of seven EC2 instances The web. Application and database sewers 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 (raffile 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: CE

NEW QUESTION 186
You are designing an SSUTLS solution that requires HTTPS clients to be authenticated by the Webserver using client certificate authentication. The solution must be resilient.
Which of the following options would you consider for configuring the web server infrastructure? (Choose 2 answers)

A. Configure ELB with TCP listeners on TCP/443. And place the Web servers behind it.
B. Configure your Web servers with EIP

C. Place the Web servers in a Route53 Record Set and configure health checks against all Web servers.
D. Configure ELB with HTTPS listeners, and place the Web sewers behind it.
E. Configure your web sewers as the origins for a CIoudFront distributio
F. Use custom SSL certificates on your C|oudFront distribution.

Answer: AB

NEW QUESTION 188
A company is building a voting system for a popular TV show, viewers win watch the performances then visit the show's website to vote for their favorite performer. It is expected that in a short period of time after the show has finished the site will receive millions of visitors. The visitors will first login to the site using their Amazon.com credentials and then submit their vote. After the voting is completed the page will display the vote totals. The company needs to build the site such that can handle the rapid influx of traffic while maintaining good performance but also wants to keep costs to a minimum. Which of the design patterns below should they use?

A. Use CIoudFront and an Elastic Load balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user then process the users vote and store the result into a multi-AZ Relational Database Service instance.
B. Use CIoudFront and the static website hosting feature of S3 with the Javascript SDK to call the Login With Amazon service to authenticate the user, use IAM Roles to gain permissions to a DynamoDB tableto store the users vote.
C. Use CIoudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login with Amazon service to authenticate the user, the web servers will process the users vote and store the result into a DynamoDB table using IAM Roles for EC2 instances to gain permissions to the DynamoDB table.
D. Use CIoudFront and an Elastic Load Balancer in front of an auto-scaled set of web servers, the web servers will first call the Login With Amazon service to authenticate the user, the web sewers win process the users vote and store the result into an SQS queue using IAM Roles for EC2 Instances to gain permissions to the SQS queu
E. A set of application sewers will then retrieve the items from the queue and store the result into a DynamoDB table.

Answer: D

NEW QUESTION 190
You are responsible for a web application that consists of an Elastic Load Balancing (ELB) load balancer in front of an Auto Scaling group of Amazon Elastic Compute Cloud (EC2) instances. For a recent deployment of a new version of the application, a new Amazon Machine Image (AMI) was created, and the Auto Scaling group was updated with a new launch configuration that refers to this new AMI. During the deployment, you received complaints from users that the website was responding with errors. All instances passed the ELB health checks.
What should you do in order to avoid errors for future deployments? (Choose 2 answer)

A. Add an Elastic Load Balancing health check to the Auto Scaling grou
B. Set a short period for the health checks to operate as soon as possible in order to prevent premature registration of the instance to theload balancer.
C. Enable EC2 instance C|oudWatch alerts to change the launch configuration's AMI to the previous on
D. Gradually terminate instances that are using the new AMI.
E. Set the Elastic Load Balancing health check configuration to target a part of the application that fully tests application health and returns an error if the tests fail.
F. Create a new launch configuration that refers to the new AMI, and associate it with the grou
G. Double the size of the group, wait for the new instances to become healthy, and reduce back to the original size.If new instances do not become healthy, associate the previous launch configuration.
H. Increase the Elastic Load Balancing Unhealthy Threshold to a higher value to prevent an unhealthy instance from going into service behind the load balancer.

Answer: CD

NEW QUESTION 191
Which is a valid Amazon Resource name (ARN) for IAM?

A. aws:iam::123456789012:instance-profile/\Nebserver
B. arn:aws:iam::123456789012:instance-profile/Webserver
C. 123456789012:aws:iam::instance-profi|e/\Nebserver
D. arn:aws:iam::123456789012::instance-profile/\Nebserver

Answer: B

NEW QUESTION 192
A company is running a batch analysis every hour on their main transactional DB, running on an RDS MySQL instance, to populate their central Data Warehouse running on Redshift. During the execution of the batch, their transactional applications are very slow. When the batch completes they need to update the top management dashboard with the new data. The dashboard is produced by another system running on-premises that is currently started when a manually-sent email notifies that an update is required. The on-premises system cannot be modified because is managed by another team.
How would you optimize this scenario to solve performance issues and automate the process as much as possible?

A. Replace RDS with Redshift for the batch analysis and SNS to notify the on-premises system to update the dashboard
B. Replace RDS with Redshift for the oaten analysis and SQS to send a message to the on-premises system to update the dashboard
C. Create an RDS Read Replica for the batch analysis and SNS to notify me on-premises system to update the dashboard
D. Create an RDS Read Replica for the batch analysis and SQS to send a message to the on-premises system to update the dashboard.

Answer: A

NEW QUESTION 197
You are running a successful multitier web application on AWS and your marketing department has asked you to add a reporting tier to the application. The reporting tier will aggregate and publish status reports every 30 minutes from user-generated information that is being stored in your web application s database. You are currently running a Multi-AZ RDS MySQL instance for the database tier. You also have implemented Elasticache as a database caching layer between the application tier and database tier. Please select the answer that will allow you to successfully implement the reporting tier with as little impact as possible to your database.

A. Continually send transaction logs from your master database to an S3 bucket and generate the reports off the S3 bucket using S3 byte range requests.
B. Generate the reports by querying the synchronously replicated standby RDS NIySQL instance maintained through Nlulti-AZ.
C. Launch a RDS Read Replica connected to your Multi AZ master database and generate reports by querying the Read Replica.
D. Generate the reports by querying the EIastiCache database caching tie

Answer: C

NEW QUESTION 200
Your company plans to host a large donation website on Amazon Web Sewices (AWS). You anticipate a large and undetermined amount of traffic that will create many database writes. To be certain that you do not drop any writes to a database hosted on AWS. Which sewice should you use?

A. Amazon RDS with provisioned IOPS up to the anticipated peak write throughput.
B. Amazon Simple Queue Service (SQS) for capturing the writes and draining the queue to write to the database.
C. Amazon EIastiCache to store the writes until the writes are committed to the database.
D. Amazon DynamoDB with provisioned write throughput up to the anticipated peak write throughpu

Answer: B

NEW QUESTION 203
You need a persistent and durable storage to trace call actMty of an IVR (Interactive Voice Response) system. Call duration is mostly in the 2-3 minutes timeframe. Each traced call can be either active or terminated. An external application needs to know each minute the list of currently active calls. Usually there are a few calls/second, but once per month there is a periodic peak up to 1000 calls/second for a few hours. The system is open 24/7 and any downtime should be avoided. Historical data is periodically archived to files. Cost saving is a priority for this project.
What database implementation would better fit this scenario, keeping costs as low as possible?

A. Use DynamoDB with a "CaIIs" table and a Global Secondary Index on a "State" attribute that can equal to "active" or "terminated". In this way the Global Secondary Index can be used for all items in the table.
B. Use RDS Multi-AZ with a "CALLS" table and an indexed "STATE" field that can be equal to "ACT|VE"or 'TERMINATED". In this way the SQL query is optimized by the use of the Index.
C. Use RDS Nlulti-AZ with two tables, one for "ACT|VE_CALLS" and one for "TERMINATED_CALLS". In this way the "ACTIVE_CALLS" table is always small and effective to access.
D. Use DynamoDB with a "CaIIs" table and a Global Secondary Index on a "Is Active" attribute that is present for active calls onl
E. In this way the Global Secondary Index is sparse and more effective.

Answer: C

NEW QUESTION 205
You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC.
The configuration is as follows: VPC: vpc-2f8bc447
IGW: igw-2d8bc445 NACL: ad-208bc448
Subnets and Route Tables: Web sewers: subnet-258bc44d
Application servers: subnet-248bc44c Database sewers: subnet-9189c6f9 Route Tables:
rrb-218bc449 rtb-238bc44b Associations:
subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b
You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.
Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these sewers to retrieve updates from the Internet?

A. Create a bastion and NAT instance in subnet-258bc44d, and add a route from rtb- 238bc44b to the NAT instance.
B. Add a route from rtb-238bc44b to igw-2d8bc445 and add a bastion and NAT instance within subnet-248bc44c.
C. Create a bastion and NAT instance in subnet-248bc44c, and add a route from rtb- 238bc44b to subneb258bc44d.
D. Create a bastion and NAT instance in subnet-258bc44d, add a route from rtb-238bc44b toIgw-2d8bc445, and a new NACL that allows access between subnet-258bc44d and subnet-248bc44

Answer: A

NEW QUESTION 207
Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

A. Use OAuth 2.0 to retrieve temporary AWS security credentials to enable your NOC members to sign in to the AWS Management Console.
B. Use web Identity Federation to retrieve AWS temporary security credentials to enable your NOC members to sign in to the AWS Management Console.
C. Use your on-premises SAML 2.0-compliant identity provider (IDP) to grant the NOC members federated access to the AWS Management Console via the AWS single sign-on (SSO) endpoint.
D. Use your on-premises SAML2.0-compliam identity provider (IDP) to retrieve temporary security credentials to enable NOC members to sign in to the AWS Management Console.

Answer: D

NEW QUESTION 208
You have an application running on an EC2 Instance which will allow users to download flies 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 user for the application with permissions that allow list access to the S3 bucket launch the instance as the IANI user and retrieve the IAM user's credentials from the EC2 instance user data.
C. Create an IAM role for EC2 that allows list access to objects in the S3 bucke
D. Launch the instance with the role, and retrieve the roIe's credentials from the EC2 Instance metadata
E. Create an IAM user for the application with permissions that allow list access to the S3 bucke
F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

Answer: C

NEW QUESTION 211
You are developing a new mobile application and are considering storing user preferences in AWS.2w This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size Additionally 5 million customers are expected to use the application on a regular basis. The solution needs to be
cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

A. Setup an RDS MySQL instance in 2 availability zones to store the user preference dat
B. Deploy apublic facing application on a server in front of the database to manage security and access credentials
C. Setup a DynamoDB table with an item for each user having the necessary attributes to hold the user preference
D. The mobile application will query the user preferences directly from the DynamoDB tabl
E. Utilize ST
F. Web Identity Federation, and DynamoDB Fine Grained Access Control to authenticate and authorize access.
G. Setup an RDS MySQL instance with multiple read replicas in 2 availability zones to store the user preference data .The mobile application will query the user preferences from the read replica
H. Leverage the MySQL user management and access prMlege system to manage security and access credentials.
I. Store the user preference data in S3 Setup a DynamoDB table with an item for each user and an item attribute pointing to the user' S3 objec
J. The mobile application will retrieve the S3 URL from DynamoDB and then access the S3 object directly utilize STS, Web identity Federation, and S3 ACLs to authenticate and authorize access.

Answer: B

NEW QUESTION 216
You deployed your company website using Elastic Beanstalk and you enabled log file rotation to S3. An Elastic Map Reduce job is periodically analyzing the logs on S3 to build a usage dashboard that you share with your CIO.
You recently improved overall performance of the website using Cloud Front for dynamic content delivery and your website as the origin.
After this architectural change, the usage dashboard shows that the traffic on your website dropped by an order of magnitude. How do you fix your usage dashboard'?

A. Enable Cloud Front to deliver access logs to S3 and use them as input of the Elastic Map Reduce job.
B. Turn on Cloud Trail and use trail log tiles on S3 as input of the Elastic Map Reduce job
C. Change your log collection process to use Cloud Watch ELB metrics as input of the Elastic MapReduce job
D. Use Elastic Beanstalk "Rebuild Environment" option to update log delivery to the Elastic lV|ap Reduce job.
E. Use Elastic Beanstalk 'Restart App server(s)" option to update log delivery to the Elastic Map Reduce job.

Answer: D

NEW QUESTION 217
Select the correct set of options. These are the initial settings for the default security group:

A. Allow no inbound traffic, Allow all outbound traffic and Allow instances associated with this security group to talk to each other
B. Allow all inbound traffic, Allow no outbound traffic and Allow instances associated with this security group to talk to each other
C. Allow no inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other
D. Allow all inbound traffic, Allow all outbound traffic and Does NOT allow instances associated with this security group to talk to each other

Answer: A

NEW QUESTION 221
How can an EBS volume that is currently attached to an EC2 instance be migrated from one Availability Zone to another?

A. Detach the volume and attach it to another EC2 instance in the other AZ.
B. Simply create a new volume in the other AZ and specify the original volume as the source.
C. Create a snapshot of the volume, and create a new volume from the snapshot in the other AZ.
D. Detach the volume, then use the ec2-migrate-voiume command to move it to another AZ.

Answer: C

NEW QUESTION 222
When you put objects in Amazon S3, what is the indication that an object was successfully stored?

A. A HTTP 200 result code and MD5 checksum, taken together, indicate that the operation was successful.
B. Amazon S3 is engineered for 99.999999999% durabilit
C. Therefore there is no need to confirm that data was inserted.
D. A success code is inserted into the S3 object metadata.
E. Each S3 account has a special bucket named _s3_log
F. Success codes are written to this bucket witha timestamp and checksum.

Answer: A

NEW QUESTION 225
A customer is deploying an SSL enabled web application to AWS and would like to implement a separation of roles between the EC2 service administrators that are entitled to login to instances as well as making API calls and the security officers who will maintain and have exclusive access to the application's X.509 certificate that contains the private key.

A. Upload the certificate on an S3 bucket owned by the security officers and accessible only by EC2 Role of the web servers.
B. Configure the web servers to retrieve the certificate upon boot from an CIoudHSM is managed by the security officers.
C. Configure system permissions on the web servers to restrict access to the certificate only to the authority security officers
D. Configure IAM policies authorizing access to the certificate store only to the security officers and terminate SSL on an ELB.

Answer: D

NEW QUESTION 227
You have recently joined a startup company building sensors to measure street noise and air quality in urban areas. The company has been running a pilot deployment of around 100 sensors for 3 months each sensor uploads 1KB of sensor data every minute to a backend hosted on AWS.
During the pilot, you measured a peak or 10 IOPS on the database, and you stored an =average of 3GB of sensor data per month in the database.
The current deployment consists of a load-balanced auto scaled Ingestion layer using EC2 instances and a PostgreSQL RDS database with 500GB standard storage.
The pilot is considered a success and your CEO has managed to get the attention or some potential investors. The business plan requires a deployment of at least 100K sensors which needs to be supported by the backend. You also need to store sensor data for at least two years to be able to compare year over year Improvements.
To secure funding, you have to make sure that the platform meets these requirements and leaves room for further scaling. Which setup win meet the requirements?

A. Add an SQS queue to the ingestion layer to buffer writes to the RDS instance
B. Ingest data into a DynamoDB table and move old data to a Redshift cluster
C. Replace the RDS instance with a 6 node Redshift cluster with 96TB of storage
D. Keep the current architecture but upgrade RDS storage to 3TB and 10K provisioned IOPS

Answer: C

NEW QUESTION 232
Your firm has uploaded a large amount of aerial image data to S3 In the past, in your on-premises environment, you used a dedicated group of servers to oaten process this data and used Rabbit MQ - An open source messaging system to get job information to the servers. Once processed the data would go to tape and be shipped offsite. Your manager told you to stay with the current design, and leverage AWS archival storage and messaging services to minimize cost. Which is correct?

A. Use SQS for passing job messages use Cloud Watch alarms to terminate EC2 worker instances when they become idl
B. Once data is processed, change the storage class of the S3 objects to Reduced Redundancy Storage.
C. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed,
D. Change the storage class of the S3 objects to Reduced Redundancy Storag
E. Setup Auto-Scaled workers triggered by queue depth that use spot instances to process messages in SOS Once data is processed, change the storage class of the S3 objects to Glacier.
F. Use SNS to pass job messages use Cloud Watch alarms to terminate spot worker instances when they become idl
G. Once data is processed, change the storage class of the S3 object to Glacier.

Answer: D

NEW QUESTION 237
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 as part of your Elastic Beanstalk definition and alter its security group to allow access to it from hosts in your application subnets.
B. 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.
C. Create your RDS instance separately and pass its DNS name to your app's DB connection string as an environment variabl
D. 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.
E. 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.

Answer: A

NEW QUESTION 240
Your application is using an ELB in front of an Auto Scaling group of web/application sewers deployed across two AZs and a MuIti-AZ RDS Instance for data persistence.
The database CPU is often above 80% usage and 90% of I/O operations on the database are reads. To improve performance you recently added a single-node Memcached EIastiCache Cluster to cache frequent DB query results. In the next weeks the overall workload is expected to grow by 30%.
Do you need to change anything in the architecture to maintain the high availability or the application with the anticipated additional load? Why?

A. Yes, you should deploy two Memcached EIastiCache Clusters in different AZs because the RDS instance will not be able to handle the load if the cache node fails.
B. No, if the cache node fails you can always get the same data from the DB without having any availability impact.
C. No, if the cache node fails the automated EIastiCache node recovery feature will prevent any availability impact.
D. Yes, you should deploy the Memcached EIastiCache Cluster with two nodes in the same AZ as the RDS DB master instance to handle the load if one cache node fails.

Answer: A

NEW QUESTION 244
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 S3, with transaction logs stored in S3 every 5 minutes.
B. Use synchronous database master-slave replication between two availability zones.
C. Take hourly DB backups to EC2 Instance store volumes with transaction logs stored In S3 every 5 minutes.
D. Take 15 minute DB backups stored In Glacier with transaction logs stored in S3 every 5 minute

Answer: A

NEW QUESTION 248
The following policy can be attached to an IAM group. It lets an IAM user in that group access a "home directory" in AWS S3 that matches their user name using the console.
{
"Version": "2012-10-17",
"Statement": [
{
"Action": ["s3:*"], "Effect": "A||ow",
"Resource": ["arn:aws:s3::zbucket-name"], "Condition":{"StringLike":{"s3:prefix":["home/${aws:username}/*"]}}
}!
{
"Action":["s3:*"], "Effect":"AI|ow",
"Resource": ["arn:aws:s3:::bucket-name/home/${aws:username}/*"]
}
}

A. True
B. False

Answer: B

NEW QUESTION 252
What does elasticity mean to AWS?

A. The ability to scale computing resources up easily, with minimal friction and down with latency.
B. The ability to scale computing resources up and down easily, with minimal friction.
C. The ability to provision cloud computing resources in expectation of future demand.
D. The ability to recover from business continuity events with minimal frictio

Answer: B

NEW QUESTION 253
Your company is storing millions of sensitive transactions across thousands of 100-GB files that must be encrypted in transit and at rest. Analysts concurrently depend on subsets of files, which can consume up to 5 TB of space, to generate simulations that can be used to steer business decisions. You are required to design an AWS solution that can cost effectively accommodate the long-term storage and in-flight subsets of data.

A. Use Amazon Simple Storage Service (S3) with server-side encryption, and run simulations on subsets in ephemeral drives on Amazon EC2.
B. Use Amazon S3 with server-side encryption, and run simulations on subsets in-memory on Amazon EC2.
C. Use HDFS on Amazon EMR, and run simulations on subsets in ephemeral drives on Amazon EC2.
D. Use HDFS on Amazon Elastic MapReduce (EMR), and run simulations on subsets in-memory on Amazon Elastic Compute Cloud (EC2).
E. Store the full data set in encrypted Amazon Elastic Block Store (EBS) volumes, and regularly capturesnapshots that can be cloned to EC2 workstation

Answer: D

NEW QUESTION 257
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AWS-Certified-Solutions-Architect-Professional Practice Exam Features:

* AWS-Certified-Solutions-Architect-Professional Questions and Answers Updated Frequently

* AWS-Certified-Solutions-Architect-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Solutions-Architect-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Solutions-Architect-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year