Amazon

Exam Questions AWS-Certified-Solutions-Architect-Professional

Amazon AWS Certified Solutions Architect Professional

NEW QUESTION 1
An IAM user is trying to perform an action on an object belonging to some other root account's bucket. Which of the below mentioned options will AWS S3 not verify?

A. The object owner has provided access to the IAM user
B. Permission provided by the parent of the IAM user on the bucket
C. Permission provided by the bucket owner to the IAM user
D. Permission provided by the parent ofthe IAM user

Answer: B

Explanation:
If the IAM user is trying to perform some action on the object belonging to another AWS user's bucket, S3 will verify whether the owner of the IAM user has given sufficient permission to him. It also verifies the policy for the bucket as well as the policy defined by the object owner.
Reference:
http://docs.aws.amazon.com/AmazonS3/latest/dev/access-control-auth-workflow-object-operation.htmI

NEW QUESTION 2
An organization is planning to host a Wordpress blog as well a joomla CMS on a single instance launched with VPC. The organization wants to have separate domains for each application and assign them using Route 53. The organization may have about ten instances each with two applications as mentioned above. While launching the instance, the organization configured two separate network interfaces (primary + ENI) and wanted to have two elastic IPs for that instance. It was suggested to use a public IP from AWS instead of an elastic IP as the number of elastic IPs is restricted. What action will you recommend to the organization?

A. I agree with the suggestion but will prefer that the organization should use separate subnets with each ENI for different public IPs.
B. I do not agree as it is required to have only an elastic IP since an instance has more than one ENI and AWS does not assign a public IP to an instance with multiple ENIs.
C. I do not agree as AWS VPC does not attach a public IP to an ENI; so the user has to use only an elastic IP only.
D. I agree with the suggestion and it is recommended to use a public IP from AWS since the organization is going to use DNS with Route 53.

Answer: B

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.
The user can attach up to two ENIs with a single instance. However, AWS cannot assign a public IP when there are two ENIs attached to a single instance. It is recommended to assign an elastic IP in this scenario. If the organization wants more than 5 E|Ps they can request AWS to increase the number.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.htmI

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
An organization is planning to host an application on the AWS VPC. The organization wants dedicated instances. However, an AWS consultant advised the organization not to use dedicated instances with VPC as the design has a few limitations. Which of the below mentioned statements is not a limitation of dedicated instances with VPC?

A. All instances launched with this VPC will always be dedicated instances and the user cannot use a default tenancy model for them.
B. It does not support the AWS RDS with a dedicated tenancy VPC.
C. The user cannot use Reserved Instances with a dedicated tenancy model.
D. The EBS volume will not be on the same tenant hardware as the EC2 instance though the user has configured dedicated tenancy.

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. Dedicated instances are Amazon EC2 instances that run in a Virtual Private Cloud (VPC) on hardware that is dedicated to a single customer. The client's dedicated instances are physically isolated at the host hardware level from instances that are not dedicated instances as well as from instances that belong to other AWS accounts.
All instances launched with the dedicated tenancy model of VPC will always be dedicated instances. Dedicated tenancy has a limitation that it may not support a few services, such as RDS. Even the EBS will not be on dedicated hardware. However the user can save some cost as well as reserve some capacity by using a Reserved Instance model with dedicated tenancy.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

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
A user is planning to host a web server as well as an app server on a single EC2 instance which is a part of the public subnet of a VPC. How can the user setup to have two separate public IPs and separate security groups for both the application as well as the web server?

A. Launch VPC with two separate subnets and make the instance a part of both the subnets.
B. Launch a VPC instance with two network interface
C. Assign a separate security group and elastic IP to them.
D. Launch a VPC instance with two network interface
E. Assign a separate security group to each and AWS will assign a separate public IP to them.
F. Launch a VPC with ELB such that it redirects requests to separate VPC instances of the public subne

Answer: B

Explanation:
If you need to host multiple websites(with different IPs) on a single EC2 instance, the following is the suggested method from AWS.
Launch a VPC instance with two network interfaces
Assign elastic IPs from VPC EIP pool to those interfaces (Because, when the user has attached more than one network interface with an instance, AWS cannot assign public IPs to them.)
Assign separate Security Groups if separate Security Groups are needed
This scenario also helps for operating network appliances, such as firewalls or load balancers that have multiple private IP addresses for each network interface.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Mu|tip|e|P.htmI

NEW QUESTION 7
You have subscribed to the AWS Business and Enterprise support plan. Your business has a backlog of problems, and you need about 20 of your IAM users to open technical support cases. How many users can open technical support cases under the AWS Business and Enterprise support plan?

A. 5 users
B. 10 users
C. Unlimited
D. 1 user

Answer: C

Explanation:
In the context of AWS support, the Business and Enterprise support plans allow an unlimited number of users to open technical support cases (supported by AWS Identity and Access Management (IAM)). Reference: https://aws.amazon.com/premiumsupport/faqs/

NEW QUESTION 8
While implementing the policy keys in AWS Direct Connect, if you use and the request comes from
an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed.

A. aws:SecureTransport
B. aws:EpochIP
C. aws:SourceIp
D. aws:CurrentTime

Answer: C

Explanation:
While implementing the policy keys in Amazon RDS, if you use aws:SourceIp and the request comes from an Amazon EC2 instance, the instance's public IP address is evaluated to determine if access is allowed. Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.htmI

NEW QUESTION 9
The Principal element of an IAM policy refers to the specific entity that should be allowed or denied permission, whereas the translates to everyone except the specified entity.

A. NotPrincipa|
B. Vendor
C. Principal
D. Action

Answer: A

Explanation:
The element NotPrincipal that is included within your IAM policy statements allows you to specify an exception to a list of principals to whom the access to a specific resource is either allowed or denied. Use the NotPrincipal element to specify an exception to a list of principals. For example, you can deny access to all principals except the one named in the NotPrincipal element.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Principal

NEW QUESTION 10
Doug has created a VPC with CIDR 10.201.0.0/16 in his AWS account. In this VPC he has created a public subnet with CIDR block 10.201.31.0/24. While launching a new EC2 from the console, he is not able to assign the private IP address 10.201.31.6 to this instance. Which is the most likely reason for this issue?

A. Private address IP 10.201.31.6 is currently assigned to another interface.
B. Private IP address 10.201.31.6 is reserved by Amazon for IP networking purposes.
C. Private IP address 10.201.31.6 is blocked via ACLs in Amazon infrastructure as a part of platform security.
D. Private IP address 10.201.31.6 is not part of the associated subnet's IP address rang

Answer: A

Explanation:
In Amazon VPC, you can assign any Private IP address to your instance as long as it is: Part of the associated subnet's IP address range
Not reserved by Amazon for IP networking purposes Not currently assigned to another interface Reference: http://aws.amazon.com/vpc/faqs/

NEW QUESTION 10
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 12
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/datapipeline/faqs/

NEW QUESTION 13
One of your AWS Data Pipeline actMties has failed consequently and has entered a hard failure state after retrying thrice. You want to try it again. Is it possible to increase the number of automatic retries to more than thrice?

A. Yes, you can increase the number of automatic retries to 6.
B. Yes, you can increase the number of automatic retries to indefinite number.
C. No, you cannot increase the number of automatic retries.
D. Yes, you can increase the number of automatic retries to 10.

Answer: D

Explanation:
In AWS Data Pipeline, an actMty fails if all of its actMty attempts return with a failed state. By default, an actMty retries three times before entering a hard failure state. You can increase the number of automatic retries to 10. However, the system does not allow indefinite retries.
Reference: https://aws.amazon.com/datapipeline/faqs/

NEW QUESTION 16
True or False: In Amazon ElastiCache replication groups of Redis, for performance tuning reasons, you can change the roles of the cache nodes within the replication group, with the primary and one of the replicas exchanging roles.

A. True, however, you get lower performance.
B. FALSE
C. TRUE
D. False, you must recreate the replication group to improve performance tunin

Answer: C

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

NEW QUESTION 19
Regarding Amazon SNS, you can send notification messages to mobile devices through any of the following supported push notification services, EXCEPT:

A. Microsoft Windows Mobile Messaging (MWMM)
B. Google Cloud Messaging for Android (GCM)
C. Amazon Device Messaging (ADM)
D. Apple Push Notification Service (APNS)

Answer: A

Explanation:
In Amazon SNS, you have the ability to send notification messages directly to apps on mobile devices. Notification messages sent to a mobile endpoint can appear in the mobile app as message alerts, badge updates, or even sound alerts. Microsoft Windows Mobile Messaging (MWMM) doesn't exist and is not supported by Amazon SNS.
Reference: http://docs.aws.amazon.com/sns/latest/dg/SNSMobilePush.htm|

NEW QUESTION 22
IAM Secure And Scalable is an organization which provides scalable and secure SAAS to its clients. They are planning to host a web server and App server on AWS VPC as separate tiers. The organization wants to implement the scalability by configuring Auto Scaling and load balancer with their app servers (middle tier) too. Which of the below mentioned options suits their requirements?

A. Since ELB is internet facing, it is recommended to setup HAProxy as the Load balancer within the VPC.
B. Create an Internet facing ELB with VPC and configure all the App servers with it.
C. The user should make ELB with EC2-CLASSIC and enable SSH with it for security.
D. Create an Internal Load balancer with VPC and register all the App sewers with i

Answer: D

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.
Reference:
http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/vpc-loadbalancer-types.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 26
An organization is setting up a backup and restore system in AWS of their in premise system. The organization needs High Availability(HA) and Disaster Recovery(DR) but is okay to have a longer recovery time to save costs. Which of the below mentioned setup options helps achieve the objective of cost saving as well as DR in the most effective way?

A. Setup pre- configured sewers and create AMIs.. Use EIP and Route 53 to quickly switch over to AWS from in premise.
B. Setup the backup data on S3 and transfer data to S3 regularly using the storage gateway.
C. Setup a small instance with AutoScaling; in case of DR start diverting all the load to AWS from on premise.
D. Replicate on premise DB to EC2 at regular intervals and setup a scenario similar to the pilot ligh

Answer: B

Explanation:
AWS has many solutions for Disaster Recovery(DR) and High Availability(HA). When the organization wants to have HA and DR but are okay to have a longer recovery time they should select the option backup and restore with S3. The data can be sent to S3 using either Direct Connect, Storage Gateway or over the internet.
The EC2 instance will pick the data from the S3 bucket when started and setup the environment. This process takes longer but is very cost effective due to the low pricing of S3. In all the other options, the EC2 instance might be running or there will be AMI storage costs.
Thus, it will be a costlier option. In this scenario the organization should plan appropriate tools to take a backup, plan the retention policy for data and setup security of the data.
Reference: http://d36cz9buwru1tt.cloudfront.net/AWS_Disaster_Recovery.pdf

NEW QUESTION 27
By default, what is the maximum number of Cache Nodes you can run in Amazon ElastiCache?

A. 20
B. 50
C. 100
D. 200

Answer: A

Explanation:
In Amazon ElastiCache, you can run a maximum of 20 Cache Nodes. Reference: http://aws.amazon.com/elasticache/faqs/

NEW QUESTION 30
What feature of the load balancing service attempts to force subsequent connections to a service to be redirected to the same node as long as it is online?

A. Node balance
B. Session retention
C. Session multiplexing
D. Session persistence

Answer: D

Explanation:
Session persistence is a feature of the load balancing service. It attempts to force subsequent connections to a service to be redirected to the same node as long as it is online.
Reference:
http://docs.rackspace.com/loadbalancers/api/v1.0/clb-devguide/content/Concepts-d1e233.html

NEW QUESTION 34
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_credentials_temp_control-access_disable-perms.h tml

NEW QUESTION 35
An organization has created 5 IAM users. The organization wants to give them the same login ID but different passwords. How can the organization achieve this?

A. The organization should create each user in a separate region so that they have their own URL to login
B. The organization should create a separate login ID but give the IAM users the same alias so that each one can login with their alias
C. It is not possible to have the same login ID for multiple IAM users of the same account
D. The organization should create various groups and add each user with the same login ID to different group
E. The user can login with their own group ID

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. Whenever the organization is creating an IAM user, there should be a unique ID for each user. It is not possible to have the same login ID for multiple users. The names of users, groups, roles, instance profiles must be alphanumeric, including the following common characters: plus (+), equal (=), comma (,), period (.), at (@), and dash (-).
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.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 40
What is the maximum length for an instance profile name in AWS IAM?

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

Answer: B

Explanation:
The maximum length for an instance profile name is 128 characters.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/LimitationsOnEntities.html

NEW QUESTION 44
An organization is planning to create a secure scalable application with AWS VPC and ELB. The organization has two instances already running and each instance has an ENI attached to it in addition to a primary network interface. The primary network interface and additional ENI both have an elastic IP attached to it.
If those instances are registered with ELB and the organization wants ELB to send data to a particular EIP of the instance, how can they achieve this?

A. The organization should ensure that the IP which is required to receive the ELB traffic is attached to a primary network interface.
B. It is not possible to attach an instance with two EN|s with ELB as it will give an IP conflict error.
C. The organization should ensure that the IP which is required to receive the ELB traffic is attached to an additional ENI.
D. It is not possible to send data to a particular IP as ELB will send to any one EI

Answer: A

Explanation:
Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances. There are two ELBs available with VPC: internet facing and internal (private) ELB. For the internet facing ELB it is required that the ELB should be in a public subnet.
When the user registers a multi-homed instance (an instance that has an Elastic Network Interface (ENI) attached) with a load balancer, the load balancer will route the traffic to the IP address of the primary network interface (eth0).
Reference: http://docs.aws.amazon.com/E|asticLoadBaIancing/latest/DeveIoperGuide/gs-ec2VPC.html

NEW QUESTION 47
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 48
A user is trying to create a PIOPS EBS volume with 3 GB size and 90 IOPS. Will AWS create the volume?

A. No, since the PIOPS and EBS size ratio is less than 30
B. Yes, since the ratio between EBS and IOPS is less than 30
C. No, the EBS size is less than 4GB
D. Yes, since PIOPS is higher than 100

Answer: C

Explanation:
A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htmI#EBSVo|umeTypes_pio ps

NEW QUESTION 50
You have been given the task to define multiple AWS Data Pipeline schedules for different actMties in the same pipeline. Which of the following would successfully accomplish this task?

A. Creating multiple pipeline definition files
B. Defining multiple pipeline definitions in your schedule objects file and associating the desired schedule to the correct actMty via its schedule field
C. Defining multiple schedule objects in your pipeline definition file and associating the desired schedule to the correct actMty via its schedule field
D. Defining multiple schedule objects in the schedule field

Answer: C

Explanation:
To define multiple schedules for different actMties in the same pipeline, in AWS Data Pipeline, you should define multiple schedule objects in your pipeline definition file and associate the desired schedule to the correct actMty via its schedule field. As an example of this, it could allow you to define a pipeline in which log files are stored in Amazon S3 each hour to drive generation of an aggregate report once a day. Reference: https://aws.amazon.com/datapipeline/faqs/

NEW QUESTION 51
In a VPC, can you modify a set of DHCP options after you create them?

A. Yes, you can modify a set of DHCP options within 48 hours after creation and there are no VPCs associated with them.
B. Yes, you can modify a set of DHCP options any time after you create them.
C. No, you can't modify a set of DHCP options after you create them.
D. Yes, you can modify a set of DHCP options within 24 hours after creatio

Answer: C

Explanation:
After you create a set of DHCP options, you can't modify them. If you want your VPC to use a different set of DHCP options, you must create a new set and associate them with your VPC. You can also set up your VPC to use no DHCP options at all.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_DHCP_Options.html

NEW QUESTION 54
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
What RAID method is used on the Cloud Block Storage back-end to implement a very high level of reliability and performance?

A. RAID 1 (Mirror)
B. RAID 5 (Blocks striped, distributed parity)
C. RAID 10 (Blocks mirrored and striped)
D. RAID 2 (Bit level striping)

Answer: C

Explanation:
Cloud Block Storage back-end storage volumes employs the RAID 10 method to provide a very high level of reliability and performance.
Reference: http://www.rackspace.com/knowledge_center/product-faq/cloud-block-storage

NEW QUESTION 57
With Amazon Elastic MapReduce (Amazon EMR) you can analyze and process vast amounts of data. The cluster is managed using an open-source framework called Hadoop.
You have set up an application to run Hadoop jobs. The application reads data from DynamoDB and generates a temporary file of 100 TBs.
The whole process runs for 30 minutes and the output of the job is stored to S3. Which of the below mentioned options is the most cost effective solution in this case?

A. Use Spot Instances to run Hadoop jobs and configure them with EBS volumes for persistent data storage.
B. Use Spot Instances to run Hadoop jobs and configure them with ephermal storage for output file storage.
C. Use an on demand instance to run Hadoop jobs and configure them with EBS volumes for persistent storage.
D. Use an on demand instance to run Hadoop jobs and configure them with ephemeral storage for output file storage.

Answer: B

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

NEW QUESTION 59
In Amazon SNS, to send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following, except:

A. Device token
B. Client ID
C. Registration ID
D. Client secret

Answer: A

Explanation:
To send push notifications to mobile devices using Amazon SNS and ADM, you need to obtain the following: Registration ID and Client secret.
Reference: http://docs.aws.amazon.com/sns/latest/dg/SNSMobi|ePushPrereq.htmI

NEW QUESTION 64
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 65
How can a user list the IAM Role configured as a part of the launch config?

A. as-describe-launch-configs --iam-profiIe
B. as-describe-launch-configs --show-long
C. as-describe-launch-configs —iam-role
D. as-describe-launch-configs —roIe

Answer: B

Explanation:
As-describe-launch-configs describes all the launch config parameters created by the AWS account in the specified region. Generally it returns values, such as Launch Config name, Instance Type and AMI ID. If the user wants additional parameters, such as the IAM Profile used in the config , he has to run command: as-describe-launch-configs --show-long

NEW QUESTION 70
Which of the following is true while using an IAM role to grant permissions to applications running on Amazon EC2 instances?

A. All applications on the instance share the same role, but different permissions.
B. All applications on the instance share multiple roles and permissions.
C. MuItipIe roles are assigned to an EC2 instance at a time.
D. Only one role can be assigned to an EC2 instance at a tim

Answer: D

Explanation:
Only one role can be assigned to an EC2 instance at a time, and all applications on the instance share the same role and permissions.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/role-usecase-ec2app.htmI

NEW QUESTION 74
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/DeveloperGuide/dp-how-tasks-scheduled.html

NEW QUESTION 75
Select the correct statement about Amazon ElastiCache.

A. It makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud.
B. It allows you to quickly deploy your cache environment only if you install software.
C. It does not integrate with other Amazon Web Services.
D. It cannot run in the Amazon Virtual Private Cloud (Amazon VPC) environmen

Answer: A

Explanation:
ElastiCache is a web service that makes it easy to set up, manage, and scale a distributed in-memory cache environment in the cloud. It provides a high-performance, scalable, and cost-effective caching solution, while removing the complexity associated with deploying and managing a distributed cache environment. With ElastiCache, you can quickly deploy your cache environment, without having to provision hardware or install software.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/WhatIs.html

NEW QUESTION 77
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/latest/UserGuide/reference_policies_elements.html

NEW QUESTION 79
Regarding Identity and Access Management (IAM), Which type of special account belonging to your application allows your code to access Google services programmatically?

A. Service account
B. Simple Key
C. OAuth
D. Code account

Answer: A

Explanation:
A service account is a special Google account that can be used by applications to access Google
services programmatically. This account belongs to your application or a virtual machine (VM), instead of to an indMdual end user. Your application uses the service account to call the Google API of a service, so that the users aren't directly involved.
A service account can have zero or more pairs of service account keys, which are used to authenticate to Google. A service account key is a public/private keypair generated by Google. Google retains the public
key, while the user is given the private key.
Reference: https://cloud.google.com/iam/docs/service-accounts

NEW QUESTION 81
What happens when Dedicated instances are launched into a VPC?

A. If you launch an instance into a VPC that has an instance tenancy of dedicated, you must manually create a Dedicated instance.
B. If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is created as a Dedicated instance, only based on the tenancy of the instance.
C. If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance isautomatically a Dedicated instance, regardless of the tenancy of the instance.

D. None of these are tru

Answer: C

Explanation:
If you launch an instance into a VPC that has an instance tenancy of dedicated, your instance is automatically a Dedicated instance, regardless of the tenancy of the instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/dedicated-instance.html

NEW QUESTION 84
An organization is setting up RDS for their applications. The organization wants to secure RDS access with VPC. Which of the following options is not required while designing the RDS with VPC?

A. The organization must create a subnet group with public and private subnet
B. Both the subnets can be in the same or separate AZ.
C. The organization should keep minimum of one IP address in each subnet reserved for RDS failover.
D. If the organization is connecting RDS from the internet it must enable the VPC attributes DNS hostnames and DNS resolution.
E. The organization must create a subnet group with VPC using more than one subnet which are a part of separate AZs.

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, such as RDS into a virtual network that the user has defined. Subnets are segments of a VPC's IP address range that the user can designate to a group of VPC resources based on security and operational needs. A DB subnet group is a collection of subnets (generally private) that the user can create in a VPC and assign to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating the DB instances.
Each DB subnet group should have subnets in at least two Availability Zones in a given region. If the RDS instance is required to be accessible from the internet the organization must enable the VPC attributes, DNS hostnames and DNS resolution. For each RDS DB instance that the user runs in a VPC, he should reserve at least one address in each subnet in the DB subnet group for use by Amazon RDS for recovery actions.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 87
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 88
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.htmI

NEW QUESTION 92
A user has created a VPC with CIDR 20.0.0.0/16. The user has created one subnet with CIDR 20.0.0.0/16 in this VPC. The user is trying to create another subnet with the same VPC for CIDR 20.0.0.1/24. What will happen in this scenario?

A. The VPC will modify the first subnet CIDR automatically to allow the second subnet IP range
B. The second subnet will be created

C. It will throw a CIDR overlaps error
D. It is not possible to create a subnet with the same CIDR as VPC

Answer: C

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. A user can create a subnet with VPC and launch instances inside that subnet. The user can create a subnet with the same size of VPC. However, he cannot create any other subnet since the CIDR of the second subnet will conflict with the first subnet.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION 95
True or False: The Amazon ElastiCache clusters are not available for use in VPC at this time.

A. TRUE
B. True, but they are available only in the GovCloud.
C. True, but they are available only on request.
D. FALSE

Answer: D

Explanation:
Amazon Elasticache clusters can be run in an Amazon VPC. With Amazon VPC, you can define a virtual network topology and customize the network configuration to closely resemble a traditional network that you might operate in your own datacenter. You can now take advantage of the manageability, availability and scalability benefits of Amazon ElastiCache Clusters in your own isolated network. The same functionality of Amazon ElastiCache, including automatic failure detection, recovery, scaling, auto discovery, Amazon CloudWatch metrics, and software patching, are now available in Amazon VPC. Reference: http://aws.amazon.com/about-aws/whats-new/2012/12/20/amazon-elasticache-announces-support-for-a mazon-vpc/

NEW QUESTION 97
In Amazon Redshift, how many slices does a dw2.8xlarge 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 100
In the context of Amazon E|astiCache CLI, which of the following commands can you use to view all ElastiCache instance events for the past 24 hours?

A. elasticache-events --duration 24
B. elasticache-events --duration 1440
C. elasticache-describe-events --duration 24
D. elasticache describe-events --source-type cache-cluster --duration 1440

Answer: D

Explanation:
In Amazon ElastiCache, the code "aws elasticache describe-events --source-type cache-cluster
--duration 1440" is used to list the cache-cluster events for the past 24 hours (1440 minutes). Reference:
http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/ECEvents.Viewing.html

NEW QUESTION 102
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 105
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 108
Which of following IAM policy elements lets you specify an exception to a list of actions?

A. NotException
B. ExceptionAction
C. Exception
D. NotAction

Answer: D

Explanation:
The NotAction element lets you specify an exception to a list of actions. Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_ElementDescriptions.html

NEW QUESTION 112
You are setting up some EBS volumes for a customer who has requested a setup which includes a RAID (redundant array of inexpensive disks). AWS has some recommendations for RAID setups. Which RAID setup is not recommended for Amazon EBS?

A. RAID 1 only
B. RAID 5 only
C. RAID 5 and RAID 6
D. RAID 0 only

Answer: C

Explanation:
With Amazon EBS, you can use any of the standard RAID configurations that you can use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating
system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together.
RAID 5 and RAID 6 are not recommended for Amazon EBS because the parity write operations of these RAID modes consume some of the IOPS available to your volumes.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/raid-config.html

NEW QUESTION 116
What is the role of the PoIIForTask action when it is called by a task runner in AWS Data Pipeline?

A. It is used to retrieve the pipeline definition.
B. It is used to report the progress of the task runner to AWS Data Pipeline.
C. It is used to receive a task to perform from AWS Data Pipeline.
D. It is used to inform AWS Data Pipeline of the outcome when the task runner completes a tas

Answer: C

Explanation:
Task runners call Po||ForTask to receive a task to perform from AWS Data Pipeline. If tasks are ready in the work queue, PoIIForTask returns a response immediately. If no tasks are available in the queue, PoIIForTask uses long-polling and holds on to a poll connection for up to 90 seconds, during which time any newly scheduled tasks are handed to the task agent. Your remote worker should not call PoIIForTask again on the same worker group until it receives a response, and this may take up to 90 seconds. Reference: http://docs.aws.amazon.com/datapipeline/latest/APIReference/AP|_Po||ForTask.htmI

NEW QUESTION 119
Who is responsible for modifying the routing tables and networking ACLs in a VPC to ensure that a DB instance is reachable from other instances in the VPC?

A. AWS administrators
B. The owner of the AWS account
C. Amazon
D. The DB engine vendor

Answer: B

Explanation:
You are in charge of configuring the routing tables of your VPC as well as the network ACLs rules needed to make your DB instances accessible from all the instances of your VPC that need to communicate with it.
Reference: http://aws.amazon.com/rds/faqs/

NEW QUESTION 124
An organization is planning to host a web application in the AWS VPC. The organization does not want to host a database in the public cloud due to statutory requirements. How can the organization setup in this scenario?

A. The organization should plan the app server on the public subnet and database in the organization's data center and connect them with the VPN gateway.
B. The organization should plan the app server on the public subnet and use RDS with the private subnet for a secure data operation.
C. The organization should use the public subnet for the app server and use RDS with a storage gateway to access as well as sync the data securely from the local data center.
D. The organization should plan the app server on the public subnet and database in a private subnet so it will not be in the public cloud.

Answer: A

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. If the user wants to connect VPC from his own data centre, he can setup a public and VPN only subnet which uses hardware VPN access to
connect with his data centre. When the user has configured this setup with Wizard, it will create a virtual private gateway to route all the traffic of the VPN subnet.
If the virtual private gateway is attached with VPC and the user deletes the VPC from the console it will first automatically detach the gateway and only then delete the VPC.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

NEW QUESTION 127
A user is trying to create a PIOPS EBS volume with 4000 IOPS and 100 GB size. AWS does not allow the user to create this volume. What is the possible root cause for this?

A. PIOPS is supported for EBS higher than 500 GB size
B. The maximum IOPS supported by EBS is 3000
C. The ratio between IOPS and the EBS volume is higher than 30
D. The ratio between IOPS and the EBS volume is lower than 50

Answer: C

Explanation:
A Provisioned IOPS (SSD) volume can range in size from 4 GiB to 16 TiB and you can provision up to 20,000 IOPS per volume. The ratio of IOPS provisioned to the volume size requested should be a maximum of 30; for example, a volume with 3000 IOPS must be atleast 100 GB.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.html#EBSVolumeTypes_pio ps

NEW QUESTION 132
You're trying to delete an SSL certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CIoudFront."
Which of the following statements is probably the reason why you are getting this error?

A. Before you can delete an SSL certificate you need to set up https on your server.
B. Before you can delete an SSL certificate, you need to set up the appropriate access level in IAM
C. Before you can delete an SSL certificate, you need to either rotate SSL certificates or revert from using a custom SSL certificate to using the default CIoudFront certificate.
D. You can't delete SSL certificates . You need to request it from AW

Answer: C

Explanation:
CIoudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, and image files, to end users.
Every CIoudFront web distribution must be associated either with the default CIoudFront certificate or with a custom SSL certificate. Before you can delete an SSL certificate, you need to either rotate SSL certificates (replace the current custom SSL certificate with another custom SSL certificate) or revert from using a custom SSL certificate to using the default CIoudFront certificate.
Reference: http://docs.aws.amazon.com/AmazonC|oudFront/latest/DeveIoperGuide/Troubleshooting.html

NEW QUESTION 134
Can Provisioned IOPS be used on RDS instances launched in a VPC?

A. Yes, they can be used only with Oracle based instances.
B. Yes, they can be used for all RDS instances.
C. No
D. Yes, they can be used only with MySQL based instance

Answer: B

Explanation:
The basic building block of Amazon RDS is the DB instance. DB instance storage comes in three types: Magnetic, General Purpose (SSD), and Provisioned IOPS (SSD). When you buy a server, you get CPU, memory, storage, and IOPS, all bundled together. With Amazon RDS, these are split apart so that you can scale them independently. So, for example, if you need more CPU, less IOPS, or more storage, you can easily allocate them.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/RDSFAQ.PIOPS.html

NEW QUESTION 135
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 138
An organization is setting up their website on AWS. The organization is working on various security measures to be performed on the AWS EC2 instances. Which of the below mentioned security mechanisms will not help the organization to avoid future data leaks and identify security weaknesses?

A. Run penetration testing on AWS with prior approval from Amazon.
B. Perform SQL injection for application testing.
C. Perform a Code Check for any memory leaks.
D. Perform a hardening test on the AWS instanc

Answer: C

Explanation:
AWS security follows the shared security model where the user is as much responsible as Amazon. Since Amazon is a public cloud it is bound to be targeted by hackers. If an organization is planning to host their application on AWS EC2, they should perform the below mentioned security checks as a measure to find any security weakness/data leaks:
Perform penetration testing as performed by attackers to find any vulnerability. The organization must take an approval from AWS before performing penetration testing
Perform hardening testing to find if there are any unnecessary ports open Perform SQL injection to find any DB security issues
The code memory checks are generally useful when the organization wants to improve the application performance.
Reference: http://aws.amazon.com/security/penetration-testing/

NEW QUESTION 139
Is there any way to own a direct connection to Amazon Web Services?

A. No, AWS only allows access from the public Internet.
B. No, you can create an encrypted tunnel to VPC, but you cannot own the connection.
C. Yes, you can via Amazon Dedicated Connection.
D. Yes, you can via AWS Direct Connec

Answer: D

Explanation:
AWS Direct Connect links your internal network to an AWS Direct Connect location over a standard 1 gigabit or 10 gigabit Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an AWS Direct Connect router. With this connection in place, you can create virtual interfaces directly to the AWS cloud (for example, to Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3)) and to Amazon Virtual Private Cloud (Amazon VPC), bypassing Internet service providers in your network path.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

NEW QUESTION 140
In Amazon ElastiCache, which of the following statements is correct?

A. When you launch an ElastiCache cluster into an Amazon VPC private subnet, every cache node is assigned a public IP address within that subnet.
B. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
C. If your AWS account supports only the EC2-VPC platform, ElastiCache will never launch your cluster in a VPC.
D. ElastiCache is not fully integrated with Amazon Virtual Private Cloud (VPC).

Answer: B

Explanation:
The VPC must allow non-dedicated EC2 instances. You cannot use ElastiCache in a VPC that is configured for dedicated instance tenancy.
Reference: http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/AmazonVPC.EC.html

NEW QUESTION 144
You are looking to migrate your Development (Dev) and Test environments to AWS. You have decided to use separate AWS accounts to host each environment. You plan to link each accounts bill to a Master AWS account using Consolidated Billing. To make sure you Keep within budget you would like to implement a way for administrators in the Master account to have access to stop, delete and/or terminate resources in both the Dev and Test accounts. Identify which option will allow you to achieve this goal.

A. Create IAM users in the Master account with full Admin permission
B. Create cross-account roles in the Dev and Test accounts that grant the Master account access to the resources in the account by inheriting permissions from

the Master account.
C. Create IAM users and a cross-account role in the Master account that grants full Admin permissions to the Dev and Test accounts.
D. Create IAM users in the Master account Create cross-account roles in the Dev and Test accounts that have full Admin permissions and grant the Master account access.
E. Link the accounts using Consolidated Billin
F. This will give IAM users in the Master account access to resources in the Dev and Test accounts

Answer: C

NEW QUESTION 146
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 149
To serve Web traffic for a popular product your chief financial officer and IT director have purchased 10 ml large heavy utilization Reserved Instances (Rls) evenly spread across two availability zones: Route 53 is used to deliver the traffic to an Elastic Load Balancer (ELB). After several months, the product grows even more popular and you need additional capacity As a result, your company purchases two C3.2xIarge medium utilization Ris You register the two c3 2xIarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xIarge instances have significant capacity that's unused Which option is the most cost effective and uses EC2 capacity most effectively?

A. Configure Autoscaling group and Launch Configuration with ELB to add up to 10 more on-demand m1 .|arge instances when triggered by Cloudwatc
B. Shut off c3.2x|arge instances.
C. Configure ELB with two c3.2xIarge instances and use on-demand Autoscaling group for up to two additional c3.2x|arge instance
D. Shut off m1 .large instances.
E. Route traffic to EC2 m1 .large and c3.2xIarge instances directly using Route 53 latency based routing and health check
F. Shut off ELB.
G. Use a separate ELB for each instance type and distribute load to ELBs with Route 53 weighted round robin.

Answer: B

NEW QUESTION 150
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 152
A read only news reporting site with a combined web and application tier and a database tier that receives large and unpredictable traffic demands must be able to respond to these traffic fluctuations automatically. What AWS services should be used meet these requirements?

A. Stateless instances for the web and application tier synchronized using EIastiCache Memcached in an autoscaimg group monitored with CIoudWatch and RDS with read replicas.
B. Stateful instances for the web and application tier in an autoscaling group monitored with CIoudWatch and RDS with read replicas.
C. Stateful instances for the web and application tier in an autoscaling group monitored with CIoudWatc
D. And multi-AZ RDS.
E. Stateless instances for the web and application tier synchronized using EIastiCache Memcached in an autoscaling group monitored with CIoudWatch and multi-AZ RDS.

Answer: A

NEW QUESTION 156
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 159
You currently operate a web application In the AWS US-East region The application runs on an
auto-scaled layer of EC2 instances and an RDS Multi-AZ database Your IT security compliance officer has tasked you to develop a reliable and durable logging solution to track changes made to your EC2.IAM And RDS resources. The solution must ensure the integrity and confidentiality of your log data. Which of these solutions would you recommend?

A. Create a new C|oudTrai| trail with one new S3 bucket to store the logs and with the global services option selected Use IAM roles S3 bucket policies and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
B. Create a new CloudTrail with one new S3 bucket to store the logs Configure SNS to send log file delivery notifications to your management system Use IAM roles and S3 bucket policies on the S3 bucket mat stores your logs.
C. Create a new CloudTrail trail with an existing S3 bucket to store the logs and with the global services option selected Use S3 ACLs and Multi Factor Authentication (MFA) Delete on the S3 bucket that stores your logs.
D. Create three new CloudTrail trails with three new S3 buckets to store the logs one for the AWS Management console, one for AWS SDKs and one for command line tools Use IAM roles and S3 bucket policies on the S3 buckets that store your logs.

Answer: A

NEW QUESTION 163
Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and US

A. The logistic software has a 3-tierarchitecture and currently uses MySQL 5.6 for data persistenc
B. Each region has deployed its own database In the HQ region you run an hourly batch process reading data from every region to compute cross-regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements'?
C. For each regional deployment, use RDS MySQL with a master in the region and a read replica in theHQ region
D. For each regional deployment, use NIySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region
E. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region
F. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region
G. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process

Answer: A

NEW QUESTION 168
A web design company currently runs several FTP servers that their 250 customers use to upload and download large graphic files They wish to move this system to AWS to make it more scalable, but they wish to maintain customer privacy and Keep costs to a minimum.
What AWS architecture would you recommend?

A. ASK their customers to use an S3 client instead of an FTP clien
B. Create a single S3 bucket Create an IAM user for each customer Put the IAM Users in a Group that has an IAM policy that permits access to sub-directories within the bucket via use of the 'username' Policy variable.
C. Create a single S3 bucket with Reduced Redundancy Storage turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket for each customer with a Bucket Policy that permits access only to that one customer.
D. Create an auto-scaling group of FTP servers with a scaling policy to automatically scale-in when minimum network traffic on the auto-scaling group is below a given threshol
E. Load a central list of ftp users from S3 as part of the user Data startup script on each Instance.
F. Create a single S3 bucket with Requester Pays turned on and ask their customers to use an S3 client instead of an FTP client Create a bucket tor each customer with a Bucket Policy that permits access only to that one customer.

Answer: A

NEW QUESTION 170
You are responsible for a legacy web application whose server environment is approaching end of life You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following limitations:
The VM's single 10GB VNIDK is almost full; Nle virtual network interface still uses the 10IV|bps driver, which leaves your 100Mbps WAN connection completely underutilized;
It is currently running on a highly customized. Windows VM within a VMware environment; You do not have me installation media;
This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

A. Use the EC2 VM Import Connector for vCenter to import the VNI into EC2.
B. Use Import/Export to import the VNI as an ESS snapshot and attach to EC2.
C. Use S3 to create a backup of the VM and restore the data into EC2.
D. Use me ec2-bundle-instance API to Import an Image of the VNI into EC2

Answer: A

NEW QUESTION 174
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 "All 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 179
Company B is launching a new game app for mobile devices. Users will log into the game using their existing social media account to streamline data capture. Company B would like to directly save player data and scoring information from the mobile app to a DynamoDS table named Score Data When a user saves their game the progress data will be stored to the Game state S3 bucket. What is the best approach for storing data to DynamoDB and S3?

A. Use an EC2 Instance that is launched with an EC2 role providing access to the Score Data DynamoDB table and the GameState S3 bucket that communicates with the mobile app via web services.
B. Use temporary security credentials that assume a role providing access to the Score Data DynamoDB table and the Game State S3 bucket using web identity federation.
C. Use Login with Amazon allowing users to sign in with an Amazon account providing the mobile app with access to the Score Data DynamoDB table and the Game State S3 bucket.
D. Use an IAM user with access credentials assigned a role providing access to the Score Data DynamoDB table and the Game State S3 bucket for distribution with the mobile app.

Answer: B

NEW QUESTION 184
Your company is getting ready to do a major public announcement of a social media site on AWS. The website is running on EC2 instances deployed across multiple Availability Zones with a Multi-AZ RDS MySQL Extra Large DB Instance. The site performs a high number of small reads and writes per second and relies on an eventual consistency model. After comprehensive tests you discover that there is read contention on RDS MySQL. Which are the best approaches to meet these requirements? (Choose 2 answers)

A. Deploy E|astiCache in-memory cache running in each availability zone
B. Implement sharding to distribute load to multiple RDS IV|ySQL instances
C. Increase the RDS MySQL Instance size and Implement provisioned IOPS
D. Add an RDS MySQL read replica in each availability zone

Answer: AC

NEW QUESTION 185
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 firewall

Answer: CEF

NEW QUESTION 189
You must architect the migration of a web application to AWS. The application consists of Linux web servers running a custom web server. You are required to save the logs generated from the application to a durable location.
What options could you select to migrate the application to AWS? (Choose 2)

A. Create an AWS Elastic Beanstalk application using the custom web server platfor
B. Specify the web server executable and the application project and source file
C. Enable log file rotation to Amazon Simple Storage Service (S3).
D. Create Dockerfile for the applicatio
E. Create an AWS OpsWorks stack consisting of a custom laye
F. Create custom recipes to install Docker and to deploy your Docker container using the Dockerfil
G. Create customer recipes to install and configure the application to publish the logs to Amazon CloudWatch Logs.
H. Create Dockerfile for the applicatio
I. Create an AWS OpsWorks stack consisting of a Docker layer that uses the Dockerfil
J. Create custom recipes to install and configure Amazon Kineses to publish the logs into Amazon CloudWatch.
K. Create a Dockerfile for the applicatio
L. Create an AWS Elastic Beanstalk application using the Docker platform and the Dockerfil
M. Enable logging the Docker configuration to automatically publish the application log
N. Enable log file rotation to Amazon S3.
O. Use VM import/Export to import a virtual machine image of the server into AWS as an AM

P. Create an Amazon Elastic Compute Cloud (EC2) instance from AMI, and install and configure the Amazon C|oudWatch Logs agen

Q. Create a new AMI from the instanc

R. Create an AWS Elastic Beanstalk application using the AMI platform and the new AMI.

Answer: AD

NEW QUESTION 191

A web company is looking to implement an external payment service into their highly available application deployed in a VPC Their application EC2 instances are behind a public lacing ELB Auto scaling is used to add additional instances as traffic increases under normal load the application runs 2 instances in the Auto Scaling group but at peak it can scale 3x in size. The application instances need to communicate with the payment service over the Internet which requires whitelisting of all public IP addresses used to communicate with it. A maximum of 4 whitelisting IP addresses are allowed at a time and can be added through an API.

How should they architect their solution?

A. Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the MAT instances.

B. Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.

C. Whitelist the ELB IP addresses and route payment requests from the Application servers through the ELB.

D. Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instances public IP address to the payment validation whitelist API.

Answer: D

NEW QUESTION 195

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 200

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 (raffle gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load unfortunately some of these new instances fail to launch.

Which of the following could be the root caused? (Choose 2 answers)

A. AWS reserves the first and the last private IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

B. The Internet Gateway (IGW) of your VPC has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches

C. The ELB has scaled-up, adding more instances to handle the traffic spike, reducing the number of available private IP addresses for new instance launches

D. AWS reserves one IP address in each subnet's CIDR block for Route53 so you do not have enough addresses left to launch all of the new EC2 instances

E. AWS reserves the first four and the last IP address in each subnet's CIDR block so you do not have enough addresses left to launch all of the new EC2 instances

Answer: CE

NEW QUESTION 204

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 CloudFront 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 CloudFront 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 CloudFront 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 CloudFront 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 205

You are designing a connectMty solution between on-premises infrastructure and Amazon VPC. Your servers on-premises will be communicating with your VPC instances. You will be establishing IPSec tunnels over the Internet You will be using VPN gateways, and terminating the IPSec tunnels on AWS supported

customer gateways.
Which of the following objectives would you achieve by implementing an IPSec tunnel as outlined above? Choose 4 answers

A. End-to-end protection of data in transit
B. End-to-end Identity authentication
C. Data encwption across the Internet
D. Protection of data in transit over the Internet
E. Peer identity authentication between VPN gateway and customer gateway
F. Data integrity protection across the Internet

Answer: CDEF

NEW QUESTION 208
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 211
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 215
An administrator is using Amazon CIoudFormation to deploy a three tier web application that consists of a web tier and application tier that will utilize Amazon DynamoDB for storage when creating the CIoudFormation template which of the following would allow the application instance access to the DynamoDB tables without exposing API credentials?

A. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and associate the Role to the application instances by referencing an instance profile.
B. Use the Parameter section in the Cloud Formation template to nave the user input Access and Secret Keys from an already created IAM user that has me permissions required to read and write from the required DynamoDB table.
C. Create an Identity and Access Management Role that has the required permissions to read and write from the required DynamoDB table and reference the Role in the instance profile property of the application instance.
D. Create an identity and Access Management user in the CIoudFormation template that has permissions to read and write from the required DynamoDB table, use the GetAtt function to retrieve the Access and secret keys and pass them to the application instance through user-data.

Answer: C

NEW QUESTION 218
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 IV|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 219
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
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 225
A web company is looking to implement an intrusion detection and prevention system into their deployed VPC. This platform should have the ability to scale to thousands of instances running inside of the VPC. How should they architect their solution to achieve these goals?

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

Answer: C

NEW QUESTION 227
You've been hired to enhance the overall security posture for a very large e-commerce site They have a well architected multi-tier application running in a VPC that uses ELBs in front of both the web and the app tier with static assets served directly from S3 They are using a combination of RDS and DynamoOB for their dynamic data and then archMng nightly into S3 for further processing with EMR They are concerned because they found questionable log entries and suspect someone is attempting to gain unauthorized access.
Which approach provides a cost effective scalable mitigation to this kind of attack?

A. Recommend that they lease space at a DirectConnect partner location and establish a 1G DirectConnect connection to their VPC they would then establish Internet connectMty into their space, filter the traffic in hardware Web Application Firewall (WAF). And then pass the traffic through the DirectConnect connection into their application running in their VPC.
B. Add previously identified hostile source IPs as an explicit INBOUND DENY NACL to the web tier subnet
C. Add a WAF tier by creating a new ELB and an AutoScaling group of EC2 Instances running ahost-based WAF They would redirect Route 53 to resolve to the new WAF tier ELB The WAF tier would their pass the traffic to the current web tier The web tier Security Groups would be updated to only allow traffic from the WAF tier Security Group
D. Remove all but TLS 1.2 from the web tier ELB and enable Advanced Protocol Filtering This will enable the ELB itself to perform WAF functionality.

Answer: C

NEW QUESTION 232
You are designing Internet connectMty for your VPC. The Web sewers must be available on the Internet. The application must have a highly available architecture. Which alternatives should you consider? (Choose 2 answers)

A. Configure a NAT instance in your VPC Create a default route via the NAT instance and associate itwith all subnets Configure a DNS A record that points to the NAT instance public IP address.
B. Configure a C|oudFront distribution and configure the origin to point to the private IP addresses of your Web sewers Configure a Route53 CNAME record to your CloudFront distribution.
C. Place all your web servers behind ELB Configure a Route53 CNMIE to point to the ELB DNS name.
D. Assign EIPs to all web sewer
E. Configure a Route53 record set with all E|Ps, with health checks and DNS failover.
F. Configure ELB with an EIP Place all your Web servers behind ELB Configure a Route53 A record that points to the EIP.

Answer: CD

NEW QUESTION 236
Your application is using an ELB in front of an Auto Scaling group of web/application sewers deployed across two AZs and a Multi-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 ElastiCache 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 ElastiCache 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 ElastiCache node recovery feature will prevent any availability impact.
D. Yes, you should deploy the Memcached ElastiCache 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 240
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 242
You are designing the network infrastructure for an application server in Amazon VPC. Users will access all application instances from the Internet, as well as from an on-premises network. The on-premises
network is connected to your VPC over an AWS Direct Connect link. How would you design routing to meet the above requirements?

A. Configure a single routing table with a default route via the Internet gatewa
B. Propagate a default route via BGP on the AWS Direct Connect customer route
C. Associate the routing table with all VPC subnets.
D. Configure a single routing table with a default route via the Internet gatewa
E. Propagate specific routes for the on-premises networks via BGP on the AWS Direct Connect customer route
F. Associate the routing table with all VPC subnets.
G. Configure a single routing table with two default routes: on to the Internet via an Internet gateway, the other to the on-premises network via the VPN gatewa
H. Use this routing table across all subnets in the VPC.
I. Configure two routing tables: on that has a default router via the Internet gateway, and other that has a default route via the VPN gatewa
J. Associate both routing tables with each VPC subnet.

Answer: A

NEW QUESTION 247
The AWS IT infrastructure that AWS provides, complies with the following IT security standards, including:

A. SOC 1/SSAE 16/ISAE 3402 (formerly SAS 70 Type II), SOC 2 and SOC 3
B. FISMA, DIACAP, and FedRA|V|P
C. PCI DSS Level 1, ISO 27001, ITAR and FIPS 140-2
D. HIPAA, Cloud Security Alliance (CSA) and Motion Picture Association of America (NIPAA)
E. All of the above

Answer: ABC

NEW QUESTION 250
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
The following are AWS Storage services? Choose 2 Answers

A. AWS Relational Database Service (AWS RDS)
B. AWS ElastiCache
C. AWS Glacier
D. AWS Import/Export

Answer: BD

NEW QUESTION 254
You have launched an EC2 instance with four (4) 500 GB EBS Provisioned IOPS volumes attached. The EC2 instance is EBS-Optimized and supports 500 Mbps throughput between EC2 and EBS. The four EBS volumes are configured as a single RAID 0 device, and each Provisioned IOPS volume is provisioned with 4,000 IOPS (4,000 16KB reads or writes), for a total of 16,000 random IOPS on the instance. The EC2 instance initially delivers the expected 16,000 IOPS random read and write performance. Sometime later, in order to increase the total random I/O performance of the instance, you
add an additional two 500 GB EBS Provisioned IOPS volumes to the RAID. Each volume is provisioned to 4,000 |OPs like the original four, for a total of 24,000 IOPS on the EC2 instance. Monitoring shows that the EC2 instance CPU utilization increased from 50% to 70%, but the total random IOPS measured at the instance level does not increase at all.
What is the problem and a valid solution?

A. The EBS-Optimized throughput limits the total IOPS that can be utilized; use an EBSOptimized instance that provides larger throughput.
B. Small block sizes cause performance degradation, limiting the I/O throughput; configure the instance device driver and filesystem to use 64KB blocks to increase throughput.
C. The standard EBS Instance root volume limits the total IOPS rate; change the instance root volume to also be a 500GB 4,000 Provisioned IOPS volume.

D. Larger storage volumes support higher Provisioned IOPS rates; increase the provisioned volume storage of each of the 6 EBS volumes to 1TB.
E. RAID 0 only scales linearly to about 4 devices; use RAID 0 with 4 EBS Provisioned IOPS volumes, but increase each Provisioned IOPS EBS volume to 6,000 IOPS.

Answer: C

NEW QUESTION 257
A newspaper organization has a on-premises application which allows the public to search its back catalogue and retrieve indMdual newspaper pages via a website written in Java They have scanned the old newspapers into JPEGs (approx 17TB) and used Optical Character Recognition (OCR) to populate a commercial search product. The hosting platform and software are now end of life and the organization wants to migrate Its archive to AWS and produce a cost efficient architecture and still be designed for availability and durability. Which is the most appropriate?

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

Answer: C

NEW QUESTION 261
......

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

100% Actual & Verified — Instant Download, Please Click
Order The AWS-Certified-Solutions-Architect-Professional Practice Test Here