AWS-Certified-Solutions-Architect-Professional Dumps

Amazon AWS Certified Solutions Architect Professional

https://www.certleader.com/AWS-Certified-Solutions-Architect-Professional-dumps.html

NEW QUESTION 1
An organization is planning to extend their data center by connecting their DC with the AWS VPC using the VPN gateway. The organization is setting up a dynamically routed VPN connection. Which of the below mentioned answers is not required to setup this configuration?

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

Answer: B

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

NEW QUESTION 2
A customer has a website which shows all the deals available across the market. The site experiences a load of 5 large EC2 instances generally. However, a week before ThanksgMng vacation they encounter a load of almost 20 large instances. The load during that period varies over the day based on the office timings. Which of the below mentioned solutions is cost effective as well as help the website achieve better performance?

A. Setup to run 10 instances during the pre-vacation period and only scale up during the office time by launching 10 more instances using the AutoScaling schedule.
B. Keep only 10 instances running and manually launch 10 instances every day during office hours.
C. During the pre-vacation period setup 20 instances to run continuously.
D. During the pre-vacation period setup a scenario where the organization has 15 instances running and 5 instances to scale up and down using Auto Scaling based on the network I/O policy.

Answer: D

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. When the organization is experiencing varying loads and the time of the load is not known but it is higher than the routine traffic it is recommended that the organization launches a few instances before hand and then setups AutoScaling with policies which scale up and down as per the EC2 metrics, such as Network I/O or CPU utilization.
If the organization keeps all 10 additional instances as a part of the AutoScaling policy sometimes during a sudden higher load it may take time to launch instances and may not give an optimal performance. This is the reason it is recommended that the organization keeps an additional 5 instances running and the next 5 instances scheduled as per the AutoScaling policy for cost effectiveness.
Reference: http://media.amazonwebservices.com/AWS_Web_Hosting_Best_Practices.pdf

NEW QUESTION 3
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.html

NEW QUESTION 4
An organization is creating a VPC for their application hosting. The organization has created two private subnets in the same AZ and created one subnet in a separate zone. The organization wants to make a
HA system with the internal ELB. Which of these statements is true with respect to an internal ELB in this scenario?

A. ELB can support only one subnet in each availability zone.
B. ELB does not allow subnet selection; instead it will automatically select all the available subnets of the VPC.

C. If the user is creating an internal ELB, he should use only private subnets.
D. ELB can support all the subnets irrespective of their zone

Answer: A

Explanation:
The Amazon Virtual Private Cloud (Amazon VPC) allows the user to define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. The user has complete control over the virtual networking environment. Within this virtual private cloud, the user can launch AWS resources, such as an ELB, and EC2 instances.
There are two ELBs available with VPC: internet facing and internal (private) ELB. For internal servers, such as App sewers the organization can create an internal load balancer in their VPC and then place back-end application instances behind the internal load balancer. The internal load balancer will route requests to the back-end application instances, which are also using private IP addresses and only accept requests from the internal load balancer.
The Internal ELB supports only one subnet in each AZ and asks the user to select a subnet while configuring internal ELB.
Reference: http://docs.aws.amazon.com/EIasticLoadBaIancing/latest/DeveIoperGuide/USVPC_creating_basic_lb.html

NEW QUESTION 5
In Amazon EIastiCache, the failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while EIastiCache provisions a replacement for the failed cache node and it get repopulated. Which of the following is a solution to reduce this potential availability impact?

A. Spread your memory and compute capacity over fewer number of cache nodes, each with smaller capacity.
B. Spread your memory and compute capacity over a larger number of cache nodes, each with smaller capacity.
C. Include fewer number of high capacity nodes.
D. Include a larger number of cache nodes, each with high capacit

Answer: B

Explanation:
In Amazon EIastiCache, the number of cache nodes in the cluster is a key factor in the availability of your cluster running Memcached. The failure of a single cache node can have an impact on the availability of your application and the load on your back-end database while EIastiCache provisions a replacement for the failed cache node and it get repopulated. You can reduce this potential availability impact by spreading your memory and compute capacity over a larger number of cache nodes, each with smaller capacity, rather than using a fewer number of high capacity nodes.
Reference: http://docs.aws.amazon.com/AmazonEIastiCache/Iatest/UserGuide/CacheNode.Memcached.htmI

NEW QUESTION 6
IV|apMySite is setting up a web application in the AWS VPC. The organization has decided to use an AWS RDS instead of using its own DB instance for HA and DR requirements.
The organization also wants to secure RDS access. How should the web application be setup with RDS?

A. Create a VPC with one public and one private subne
B. Launch an application instance in the public subnet while RDS is launched in the private subnet.
C. Setup a public and two private subnets in different AZs within a VPC and create a subnet grou
D. Launch RDS with that subnet group.
E. Create a network interface and attach two subnets to i
F. Attach that network interface with RDS while launching a DB instance.
G. Create two separate VPCs and launch a Web app in one VPC and RDS in a separate VPC and connect them with VPC peering.

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, 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 the security and operational needs.
A DB subnet group is a collection of subnets (generally private) that a 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.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 7
When does an AWS Data Pipeline terminate the AWS Data Pipeline-managed compute resources?

A. AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 2 hours.
B. When the final actMty that uses the resources is running
C. AWS Data Pipeline terminates AWS Data Pipeline-managed compute resources every 12 hours.
D. When the final actMty that uses the resources has completed successfully orfailed

Answer: D

Explanation:
Compute resources will be provisioned by AWS Data Pipeline when the first actMty for a scheduled time that uses those resources is ready to run, and those instances will be terminated when the final actMty that uses the resources has completed successfully or failed.
Reference: https://aws.amazon.com/datapipe|ine/faqs/

NEW QUESTION 8
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 9
The Statement element, of an AWS IAM policy, contains an array of indMdual statements. Each indMdual statement is a(n) block enclosed in braces { }.

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

Answer: C

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

NEW QUESTION 10
An organization is hosting a scalable web application using AWS. The organization has configured ELB and Auto Scaling to make the application scalable. Which of the below mentioned statements is not required to be followed for ELB when the application is planning to host a web application on VPC?

A. The ELB and all the instances should be in the same subnet.
B. Configure the security group rules and network ACLs to allow traffic to be routed between the subnets in the VPC.
C. The internet facing ELB should have a route table associated with the internet gateway.
D. The internet facing ELB should be only in a public subne

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. After the user creates the public subnet, he should ensure to associate the route table of the public subnet with the internet gateway to enable the load balancer in the subnet to connect with the internet. The ELB and instances can be in a separate subnet. However, to allow communication between the instance and the
ELB the user must configure the security group rules and network ACLs to allow traffic to be routed between the subnets in his VPC.
Reference: http://docs.aws.amazon.com/EIasticLoadBaIancing/latest/DeveIoperGuide/CreateVPCForELB.htmI

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

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

Answer: B

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

NEW QUESTION 13
The two policies that you attach to an IAM role are the access policy and the trust policy. The trust policy identifies who can assume the role and grants the permission in the AWS Lambda account principal by adding the action.

A. aws:AssumeAdmin
B. Iambda:InvokeAsync
C. sts:InvokeAsync
D. sts:AssumeRole

Answer: D

Explanation:
The two policies that you attach to an IAM role are the access policy and the trust policy.
Remember that adding an account to the trust policy of a role is only half of establishing the trust relationship. By default, no users in the trusted accounts can assume the role until the administrator for that account grants the users the permission to assume the role by adding the Amazon Resource Name (ARN) of the role to an Allow element for the sts:AssumeRole action.

Reference: http://docs.aws.amazon.com/|AM/Iatest/UserGuide/id_ro|es_manage_modify.html

NEW QUESTION 18
How much memory does the cr1.8xlarge instance type provide?

A. 224 GB
B. 124 GB
C. 184 GB
D. 244 GB

Answer: D

Explanation:
The CR1 instances are part of the memory optimized instances. They offer lowest cost per GB RAM among all the AWS instance families. CR1 instances are part of the new generation of memory optimized instances, which can offer up to 244 GB RAM and run on faster CPUs (Intel Xeon E5-2670 with NUMA support) in comparison to the NI2 instances of the same family. They support cluster networking for bandwidth intensive applications. cr1.8x|arge is one of the largest instance types of the CR1 family, which can offer 244 GB RAM.
Reference: http://aws.amazon.com/ec2/instance-types/

NEW QUESTION 22
How many cg1.4xlarge on-demand instances can a user run in one region without taking any limit increase approval from AWS?

A. 20
B. 2
C. 5
D. 10

Answer: B

Explanation:
Generally AWS EC2 allows running 20 on-demand instances and 100 spot instances at a time. This limit can be increased by requesting at https://aws.amazon.com/contact-us/ec2-request. Excluding certain types of instances, the limit is lower than mentioned above. For cg1.4xlarge, the user can run only 2
on-demand instances at a time.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.htmI#|imits_ec2

NEW QUESTION 26
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/SNSMobiIePush.htm|

NEW QUESTION 30
You want to define permissions for a role in an IAM policy. Which of the following configuration formats should you use?

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

Answer: C

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

NEW QUESTION 32
Which of the following is NOT an advantage of using AWS Direct Connect?

A. AWS Direct Connect provides users access to public and private resources by using two different connections while maintaining network separation between the public and private environments.
B. AWS Direct Connect provides a more consistent network experience than Internet-based connections.
C. AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS.
D. AWS Direct Connect reduces your network cost

Answer: A

Explanation:
AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectMty between AWS and your datacenter, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.
By using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2
instances running within an Amazon Virtual Private Cloud (VPC) using private IP space, while maintaining network separation between the public and private environments.
Reference: http://aws.amazon.com/directconnect/#details

NEW QUESTION 36
Does Amazon RDS API provide actions to modify DB instances inside a VPC and associate them with DB Security Groups?

A. Yes, Amazon does this but only for MySQL RDS.
B. Yes
C. No
D. Yes, Amazon does this but only for Oracle RD

Answer: B

Explanation:
You can use the action Modify DB Instance, available in the Amazon RDS API, to pass values for the parameters DB Instance Identifier and DB Security Groups specifying the instance ID and the DB Security Groups you want your instance to be part of.
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/APIReference/API_|VIodifyDBInstance.htmI

NEW QUESTION 40
Does an AWS Direct Connect location provide access to Amazon Web Services in the region it is associated with as well as access to other US regions?

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

Answer: C

Explanation:
An AWS Direct Connect location provides access to Amazon Web Services in the region it is associated with, as well as access to other US regions. For example, you can provision a single connection to any AWS Direct Connect location in the US and use it to access public AWS services in all US Regions and AWS GovCIoud (US).
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/Welcome.html

NEW QUESTION 45
Which of the following components of AWS Data Pipeline specifies the business logic of your data management?

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

Answer: B

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

NEW QUESTION 47
In IAM, which of the following is true of temporary security credentials?

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

Answer: A

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

NEW QUESTION 52
The CFO of a company wants to allow one of his employees to view only the AWS usage report page. Which of the below mentioned IAM policy statements allows the user to have access to the AWS usage report page?

A. "Effect": "AIIow", "Action": ["Describe"], "Resource": "BiIIing"
B. "Effect": "AIIow", "Action": ["aws-portal: ViewBi||ing"], "Resource": "*"

C. "Effect": "AlIow", "Action": ["aws-portal:ViewUsage"], "Resource": "*"
D. "Effect": "AlIow", "Action": ["AccountUsage], "Resource": "*"

Answer: C

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. If the CFO wants to allow only AWS usage report page access, the policy for that IAM user will be as given below:
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "AlIow", "Action": [
"aws-portaI:ViewUsage"
]!
"Resource": "*"
} I
}
Reference: http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-permissions-ref.html

NEW QUESTION 55
In Amazon VPC, what is the default maximum number of BGP advertised routes allowed per route table?

A. 15
B. 100
C. 5
D. 10

Answer: B

Explanation:
The maximum number of BGP advertised routes allowed per route table is 100.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Appendix_Limits.html

NEW QUESTION 58
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/Iatest/UserGuide/Using_SettingUpUser.htmI

NEW QUESTION 62
A user is planning to use EBS for his DB requirement. The user already has an EC2 instance running in the VPC private subnet. How can the user attach the EBS volume to a running instance?

A. The user can create EBS in the same zone as the subnet of instance and attach that EBS to instance.
B. It is not possible to attach an EBS to an instance running in VPC until the instance is stopped.
C. The user can specify the same subnet while creating EBS and then attach it to a running instance.
D. The user must create EBS within the same VPC and then attach it to a running instance.

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. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. The instance launched will always be in the same availability zone of the respective subnet. When creating an EBS the user cannot specify the subnet or VPC. However, the user must create the EBS in the same zone as the instance so that it can attach the EBS volume to the running instance.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.htmI#VPCSubnet

NEW QUESTION 64
ExamKilIer has three separate departments and each department has their own AWS accounts. The HR department has created a file sharing site where all the on roll employees' data is uploaded. The Admin department uploads data about the employee presence in the office to their DB hosted in the VPC. The Finance department needs to access data from the HR department to know the on roll employees to calculate the salary based on the number of days that an employee is present in the office.
How can ExamKil|er setup this scenario?

A. It is not possible to configure VPC peering since each department has a separate AWS account.
B. Setup VPC peering for the VPCs of Admin and Finance.

C. Setup VPC peering for the VPCs of Finance and HR as well as between the VPCs of Finance and Admin.
D. Setup VPC peering for the VPCs of Admin and HR

Answer: C

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. 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/PeeringGuide/peering-configurations-full-access.html#t hree-vpcs-full-access

NEW QUESTION 66
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/EBSVolumeTypes.html#EBSVolumeTypes_pio ps

NEW QUESTION 67
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 70
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 75
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as long as hours.

A. 24
B. 36
C. 10
D. 48

Answer: B

Explanation:
By default, temporary security credentials for an IAM user are valid for a maximum of 12 hours, but you can request a duration as short as 15 minutes or as long as 36 hours.
Reference: http://docs.aws.amazon.com/STS/latest/UsingSTS/CreatingSessionTokens.html

NEW QUESTION 76
How does in-memory caching improve the performance of applications in ElastiCache?

A. It improves application performance by deleting the requests that do not contain frequently accessed data.
B. It improves application performance by implementing good database indexing strategies.

C. It improves application performance by using a part of instance RAM for caching important data.
D. It improves application performance by storing critical pieces of data in memory for low-latency acces

Answer: D

Explanation:
In Amazon ElastiCache, in-memory caching improves application performance by storing critical pieces of data in memory for low-latency access. Cached information may include the results of I/O-intensive database queries or the results of computationally intensive calculations.
Reference: http://aws.amazon.com/elasticache/faqs/#g4

NEW QUESTION 79
How can a user list the IAM Role configured as a part of the launch config?

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

Answer: 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 84
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. Multiple 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.html

NEW QUESTION 87
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 92
Which of the following cannot be done using AWS Data Pipeline?

A. Create complex data processing workloads that are fault tolerant, repeatable, and highly available.
B. Regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS service.
C. Generate reports over data that has been stored.
D. Move data between different AWS compute and storage services as well as on-premise data sources at specified intervals.

Answer: C

Explanation:
AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services as well as on-premise data sources at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to another AWS.
AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. AWS Data Pipeline also allows you to move and process data that was
previously locked up in on-premise data silos. Reference: http://aws.amazon.com/datapipe|ine/

NEW QUESTION 93

You create a VPN connection, and your VPN device supports Border Gateway Protocol (BGP). Which of the following should be specified to configure the VPN connection?

A. Classless routing
B. Classfull routing
C. Dynamic routing
D. Static routing

Answer: C

Explanation:
If you create a VPN connection, you must specify the type of routing that you plan to use, which will depend upon on the make and model of your VPN devices. If your VPN device supports Border Gateway Protocol (BGP), you need to specify dynamic routing when you configure your VPN connection. If your device does not support BGP, you should specify static routing.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.htmI

NEW QUESTION 98
In the context of Amazon E|astiCache CLI, which of the following commands can you use to view all EIastiCache instance events for the past 24 hours?

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

Answer: D

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

NEW QUESTION 99
In Amazon Cognito what is a silent push notification?

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

Answer: A

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

NEW QUESTION 103
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 105
Once the user has set EIastiCache for an application and it is up and running, which services, does Amazon not provide for the user:

A. The ability for client programs to automatically identify all of the nodes in a cache cluster, and to initiate and maintain connections to all of these nodes
B. Automating common administrative tasks such as failure detection and recovery, and software patching
C. Providing default Time To Live (TTL) in the AWS Elasticache Redis Implementation for different type of data.
D. Providing detailed monitoring metrics associated with your Cache Nodes, enabling you to diagnose and react to issues very quickly

Answer: C

Explanation:
Amazon provides failure detection and recovery, and software patching and monitoring tools which is called CloudWatch. In addition it provides also Auto
Discovery to automatically identify and initialize all nodes of cache cluster for Amazon EIastiCache.
Reference: http://docs.aws.amazon.com/AmazonEIastiCache/Iatest/UserGuide/WhatIs.html

NEW QUESTION 107
A user is planning to host a Highly Available system on the AWS VPC. Which of the below mentioned statements is helpful in this scenario?

A. Create VPC subnets in two separate availability zones and launch instances in different subnets.
B. Create VPC with only one public subnet and launch instances in different AZs using that subnet.
C. Create two VPCs in two separate zones and setup failover with ELB such that if one VPC fails it will divert traffic to another VPC.
D. Create VPC with only one private subnet and launch instances in different AZs using that subne

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. The VPC is always specific to a region. The user can create a VPC which can span multiple Availability Zones by adding one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span across zones.
Reference: http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.htmI#VPCSubnet

NEW QUESTION 111
A user is creating a PIOPS volume. What is the maximum ratio the user should configure between PIOPS and the volume size?

A. 5
B. 10
C. 20
D. 30

Answer: D

Explanation:
Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads that are sensitive to storage performance and consistency in random access I/O throughput. A provisioned IOPS volume can range in size from 10 GB to 1 TB and the user can provision up to 4000 IOPS per volume.
The ratio of IOPS provisioned to the volume size requested can 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.htm|

NEW QUESTION 114
To get started using AWS Direct Connect, in which of the following steps do you configure Border Gateway Protocol (BGP)?

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

Answer: C

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

NEW QUESTION 118
Which of the following components of AWS Data Pipeline polls for tasks and then performs those tasks?

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

Answer: B

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

NEW QUESTION 122
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 124
An organization has setup RDS with VPC. The organization wants RDS to be accessible from the internet. Which of the below mentioned configurations is not required in this scenario?

A. The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
B. The organization must allow access from the internet in the RDS VPC security group,
C. The organization must setup RDS with the subnet group which has an external IP.
D. The organization must enable the VPC attributes DNS hostnames and DNS resolutio

Answer: C

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources, 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 which the user assigns to the RDS DB instances. A DB subnet group allows the user to specify a particular VPC when creating DB instances. If the RDS instance is required to be accessible from the internet:
The organization must setup that the RDS instance is enabled with the VPC attributes, DNS hostnames and DNS resolution.
The organization must enable the parameter in the console which makes the RDS instance publicly accessible.
The organization must allow access from the internet in the RDS VPC security group. Reference:
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.html

NEW QUESTION 127
An organization, which has the AWS account ID as Q99988887777, has created 50 IAM users. All the users are added to the same group examkiller. If the organization has enabled that each IAM user can login with the AWS console, which AWS login URL will the IAM users use??

A. https://Q99988887777.aws.amazon.com/examkiIIer/
B. https://signin.aws.amazon.com/examki|Ier/
C. https://examkiller.signin.aws.amazon.com/999988887777/consoIe/
D. https://999988887777.signin.aws.amazon.com/consoIe/

Answer: D

Explanation:
AWS Identity and Access Management is a web service which allows organizations to manage users and user permissions for various AWS services. Once the organization has created the IAM users, they will have a separate AWS console URL to login to the AWS console. The console login URL for the IAM user will be https:// AWS_Account_ID.signin.aws.amazon.com/console/. It uses only the AWS account ID and does not depend on the group or user ID.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccountAlias.html

NEW QUESTION 131
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 136
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 C3.2xlarge medium utilization Ris You register the two c3 2xlarge instances with your ELB and quickly find that the ml large instances are at 100% of capacity and the c3 2xlarge 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 c3.2xlarge instances and use on-demand Autoscaling group for up to two additional c3.2x|arge instance
D. Shut off m1 .Iarge instances.
E. Route traffic to EC2 m1 .Iarge and c3.2xlarge 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 140
You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.exampIe.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. Dunning a DR test you notice that when you disable all web sewers in one of the regions Route53 does not automatically direct all users to the other region. What could be happening? (Choose 2 answers)

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

Answer: BE

NEW QUESTION 145
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 146
You are the new IT architect in a company that operates a mobile sleep tracking application.
When activated at night, the mobile app is sending collected data points of 1 kilobyte every 5 minutes to your backend.
The backend takes care of authenticating the user and writing the data points into an Amazon DynamoDB table.
Every morning, you scan the table to extract and aggregate last night's data on a per user basis, and store the results in Amazon S3. Users are notified via Amazon SNS mobile push notifications that new data is available, which is parsed and visualized by the mobile app.
Currently you have around 100k users who are mostly based out of North America. You have been tasked to optimize the architecture of the backend system to lower cost. What would you recommend? Choose 2 answers

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

Answer: AD

NEW QUESTION 151
You require the ability to analyze a large amount of data, which is stored on Amazon S3 using Amazon Elastic Map Reduce. You are using the cc2 8x large Instance type, whose CPUs are mostly idle during processing. Which of the below would be the most cost efficient way to reduce the runtime of the job?

A. Create more smaller flies on Amazon S3.
B. Add additional cc2 8x large instances by introducing a task group.
C. Use smaller instances that have higher aggregate I/O performance.
D. Create fewer, larger files on Amazon S3.

Answer: C

NEW QUESTION 152
Your company runs a customer facing event registration site This site is built with a 3-tier architecture with web and application tier servers and a MySQL database The application requires 6 web tier sewers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single NIySQL database. When deploying this application in a region with three availability zones (AZs) which architecture provides high availability?

A. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer), and an application tier deployed across 2 AZs with 3 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (RelationalDatabase Service) instance deployed with read replicas in the other AZ.
B. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and one RDS (Relational Database Service) Instance deployed with read replicas in the two other AZs.
C. A web tier deployed across 2 AZs with 3 EC2 (Elastic Compute Cloud) instances in each AZ inside an Auto Scaling Group behind an ELB (elastic load balancer) and an application tier deployed across 2 AZs with 3 EC2 instances m each AZ inside an Auto Scaling Group behind an ELS and a Multi-AZ RDS (Relational Database Service) deployment.
D. A web tier deployed across 3 AZs with 2 EC2 (Elastic Compute Cloud) instances in each AZ Inside an Auto Scaling Group behind an ELB (elastic load balancer). And an application tier deployed across 3 AZs with 2 EC2 instances in each AZ inside an Auto Scaling Group behind an ELB and a Multi-AZ RDS (Relational Database services) deployment.

Answer: D

NEW QUESTION 156
You would like to create a mirror image of your production environment in another region for disaster recovery purposes. Which of the following AWS resources do not need to be recreated in the second region? (Choose 2 answers)

A. Route 53 Record Sets
B. IAM Roles
C. Elastic IP Addresses (EIP)
D. EC2 Key Pairs
E. Launch configurations
F. Security Groups

Answer: AC

NEW QUESTION 157
An enterprise wants to use a third-party SaaS application. The SaaS application needs to have access to issue several API commands to discover Amazon EC2 resources running within the enterprise's account The enterprise has internal security policies that require any outside access to their environment must conform to the principles of least prMlege and there must be controls in place to ensure that the credentials used by the SaaS vendor cannot be used by any other third party. Which of the following would meet all of these conditions?

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

Answer: C

NEW QUESTION 160
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 164
Your company produces customer commissioned one-of-a-kind skiing helmets combining nigh fashion with custom technical enhancements Customers can show off their IndMduality on the ski slopes and have access to head-up-displays. GPS rear-view cams and any other technical innovation they wish to embed in the helmet.
The current manufacturing process is data rich and complex including assessments to ensure that the custom electronics and materials used to assemble the helmets are to the highest standards Assessments are a mixture of human and automated assessments you need to add a new set of assessment to model the failure modes of the custom electronics using GPUs with CUDA, across a cluster of servers with low latency networking.
What architecture would allow you to automate the existing process using a hybrid approach and ensure that the architecture can support the evolution of processes over time?

A. Use AWS Data Pipeline to manage movement of data & meta-data and assessments Use an auto-scaling group of G2 instances in a placement group.
B. Use Amazon Simple Workflow (SWF) to manages assessments, movement of data & meta-data Use an auto-scaling group of G2 instances in a placement group.
C. Use Amazon Simple Workflow (SWF) to manages assessments movement of data & meta-data Use an auto-scaling group of C3 instances with SR-IOV (Single Root I/O Virtualization).
D. Use AWS data Pipeline to manage movement of data & meta-data and assessments use auto-scaling group of C3 with SR-IOV (Single Root I/O virtualization).

Answer: B

NEW QUESTION 165
You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.exampl|e.com) and has a 2-tier architecture, with multiple application sewers and a database sewer. Remote clients use TCP to connect to the application servers. The application servers need to know the IP address of the clients in order to function properly and are currently taking that information from the TCP socket. A Multi-AZ RDS MySQL instance will be used for the database. During the migration you can change the application code, but you have to file a change request.
How would you implement the architecture on AWS in order to maximize scalability and high availability?

A. File a change request to implement Alias Resource support in the applicatio
B. Use Route 53 Alias Resource Record to distribute load on two application servers in different Azs.
C. File a change request to implement Latency Based Routing support in the applicatio
D. Use Route 53 with Latency Based Routing enabled to distribute load on two application servers in different Azs.
E. File a change request to implement Cross-Zone support in the applicatio
F. Use an ELB with a TCP Listener and Cross-Zone Load Balancing enabled, two application servers in different AZs.
G. File a change request to implement Proxy Protocol support in the applicatio
H. Use an ELB with a TCP Listener and Proxy Protocol enabled to distribute load on two application servers in different Azs.

Answer: D

NEW QUESTION 167
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 169
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 174
Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design tor the application that leverages multiple regions tor the most recently accessed content and latency sensitive portions of the wet) site The most latency sensitive component of the application involves reading user preferences to support web site personalization and ad selection. In addition to running your application in multiple regions, which option will support this application's requirements?

A. Serve user content from S3. CloudFront and use Route53 latency-based routing between ELBs in each region Retrieve user preferences from a local DynamoDB table in each region and leverage SQS to capture changes to user preferences with SOS workers for propagating updates to each table.
B. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3. C|oudFront with dynamic content and an ELB in each region Retrieve user preferences from an ElasticCache cluster in each region and leverage SNS notifications to propagate user preference changes to a worker node in each region.
C. Use the S3 Copy API to copy recently accessed content to multiple regions and serve user content from S3 CloudFront and Route53 latency-based routing Between ELBs In each region Retrieve user preferences from a DynamoDB table and leverage SQS to capture changes to user preferences with SOS workers for propagating DynamoDB updates.
D. Serve user content from S3. C|oudFront with dynamic content, and an ELB in each region Retrieve user preferences from an ElastiCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized OB to each ElastiCache cluster.

Answer: A

NEW QUESTION 179
You've been brought in as solutions architect to assist an enterprise customer with their migration of an e-commerce platform to Amazon Virtual Private Cloud (VPC) The previous architect has already deployed a 3-tier VPC.
The configuration is as follows: VPC: vpc-2f8bc447
IGW: igw-2d8bc445 NACL: ad-208bc448
Subnets and Route Tables: Web sewers: subnet-258bc44d
Application servers: subnet-248bc44c Database sewers: subnet-9189c6f9 Route Tables:
rrb-218bc449 rtb-238bc44b Associations:
subnet-258bc44d : rtb-218bc449 subnet-248bc44c : rtb-238bc44b subnet-9189c6f9 : rtb-238bc44b
You are now ready to begin deploying EC2 instances into the VPC Web servers must have direct access to the internet Application and database servers cannot have direct access to the internet.
Which configuration below will allow you the ability to remotely administer your application and database servers, as well as allow these sewers to retrieve updates from the Internet?

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

Answer: A

NEW QUESTION 181
You are designing a multi-platform web application for AWS The application will run on EC2 instances and will be accessed from PCs. tablets and smart phones Supported accessing platforms are Windows, MacOS, IOS and Android Separate sticky session and SSL certificate setups are required for different platform types which of the following describes the most cost effective and performance efficient architecture setup?

A. Setup a hybrid architecture to handle session state and SSL certificates on-prem and separate EC2 Instance groups running web applications for different platform types running in a VPC.
B. Set up one ELB for all platforms to distribute load among multiple instance under it Each EC2 instance implements ail functionality for a particular platform.
C. Set up two ELBs The first ELB handles SSL certificates for all platforms and the second ELB handles session stickiness for all platforms for each ELB run separate EC2 instance groups to handle the web application for each platform.

D. Assign multiple ELBS to an EC2 instance or group of EC2 instances running the common components of the web application, one ELB for each platform type Session stickiness and SSL termination are done at the ELBs.

Answer: D

NEW QUESTION 183
An administrator is using Amazon CloudFormation 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 CloudFormation 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 CloudFormation 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 187
You have an application running on an EC2 Instance which will allow users to download flies from a private S3 bucket using a pre-signed URL. Before generating the URL the application should verify the existence of the file in S3.
How should the application use AWS credentials to access the S3 bucket securely?

A. Use the AWS account access Keys the application retrieves the credentials from the source code of the application.
B. Create an IAM user for the application with permissions that allow list access to the S3 bucket launch the instance as the IANI user and retrieve the IAM user's credentials from the EC2 instance user data.
C. Create an IAM role for EC2 that allows list access to objects in the S3 bucke
D. Launch the instance with the role, and retrieve the role's credentials from the EC2 Instance metadata
E. Create an IAM user for the application with permissions that allow list access to the S3 bucke
F. The application retrieves the IAM user credentials from a temporary directory with permissions that allow read access only to the application user.

Answer: C

NEW QUESTION 191
Your application provides data transformation services. Files containing data to be transformed are first uploaded to Amazon S3 and then transformed by a fileet of spot EC2 instances. Files submitted by your premium customers must be transformed with the highest priority. How should you implement such a system?

A. Use a DynamoDB table with an attribute defining the priority leve
B. Transformation instances will scan the table for tasks, sorting the results by priority level.
C. Use Route 53 latency based-routing to send high priority tasks to the closest transformation instances.
D. Use two SQS queues, one for high priority messages, the other for default priorit
E. Transformation instances first poll the high priority queue; if there is no message, they poll the default priority queue.
F. Use a single SQS queu
G. Each message contains the priority leve
H. Transformation instances poll high-priority messages first.

Answer: C

NEW QUESTION 195
Which of the following are characteristics of Amazon VPC subnets? Choose 2 answers

A. Each subnet spans at least 2 Availability Zones to provide a high-availability environment.
B. Each subnet maps to a single Availability Zone.
C. CIDR block mask of /25 is the smallest range supported.
D. By default, all subnets can route between each other, whether they are private or public.
E. Instances in a private subnet can communicate with the Internet only if they have an Elastic I

Answer: AE

NEW QUESTION 200
A company is storing data on Amazon Simple Storage Service (S3). The company's security policy mandates that data is encrypted at rest. Which of the following methods can achieve this?
Choose 3 answers

A. Use Amazon S3 server-side encryption with AWS Key Management Service managed keys.
B. Use Amazon S3 server-side encryption with customer-provided keys.
C. Use Amazon S3 server-side encryption with EC2 key pair.
D. Use Amazon S3 bucket policies to restrict access to the data at rest.
E. Encrypt the data on the client-side before ingesting to Amazon S3 using their own master key.
F. Use SSL to encrypt the data while in transit to Amazon S3.

Answer: ABE

NEW QUESTION 204

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

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

Answer: D

NEW QUESTION 207
You are implementing AWS Direct Connect. You intend to use AWS public service end points such as Amazon S3, across the AWS Direct Connect link. You want other Internet traffic to use your existing link to an Internet Service Provider.
What is the correct way to configure AWS Direct connect for access to services such as Amazon S3?

A. Configure a public Interface on your AWS Direct Connect link Configure a static route via your AWS Direct Connect link that points to Amazon S3 Advertise a default route to AWS using BGP.
B. Create a private interface on your AWS Direct Connect lin
C. Configure a static route via your AWS Direct connect link that points to Amazon S3 Configure specific routes to your network in your VPC.
D. Create a public interface on your AWS Direct Connect link Redistribute BGP routes into your existing routing infrastructure; advertise specific routes for your network to AWS.
E. Create a private interface on your AWS Direct connect lin
F. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Answer: C

NEW QUESTION 209
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 CIoudFormation 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 CIoudSearch 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 CIoudFront 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 214
......

Thank You for Trying Our Product

* 100% Pass or Money Back

    All our products come with a 90-day Money Back Guarantee.

* One year free update

    You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

    We currently serve more than 30,000,000 customers.

* Shop Securely

    All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-Solutions-Architect-Professional Exam with Our Prep Materials Via below:

https://www.certleader.com/AWS-Certified-Solutions-Architect-Professional-dumps.html