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
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 2
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 3
How many g2.2xlarge 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: C

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 g2.2xlarge, the user can run only 5
on-demand instance at a time.
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_|imits.html#limits_ec2

NEW QUESTION 4
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 5
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 NotPrincipa| 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 NotPrincipa| element.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Principal

NEW QUESTION 6
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 7
A user is configuring MySQL RDS with PIOPS. What should be the minimum size of DB storage provided by the user?

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

Answer: D

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

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

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

Answer: C

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

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

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

Answer: C

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

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

A. You can use multiple compute resources on the same pipeline by defining multiple cluster objects in your definition file and associating the cluster to use for each actMty via its runsOn field.
B. You can use multiple compute resources on the same pipeline by defining multiple cluster definition files.

C. You can use multiple compute resources on the same pipeline by defining multiple clusters for your actMty.
D. You cannot use multiple compute resources on the same pipelin

Answer: A

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

NEW QUESTION 10
The two policies that you attach to an IAM role are the access policy and the trust policy. The trust policy identifies who can assume the role and grants the permission in the AWS Lambda account principal by adding the action.

A. aws:AssumeAdmin
B. Iambda:InvokeAsync
C. sts:|nvokeAsync
D. sts: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 15
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/EIasticLoadBalancing/latest/DeveIoperGuide/vpc-Ioadbalancer-types.html

NEW QUESTION 20
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 24
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 25
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 30
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/e|asticache/faqs/

NEW QUESTION 32
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 35
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. EC2 bandwidth
C. EBS volume size
D. Instance type is not EBS optimized

Answer: C

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 37
A bucket owner has allowed another account's IAM users to upload or access objects in his bucket. The IAM user of Account A is trying to access an object created by the IAM user of account B. What will happen in this scenario?

A. It is not possible to give permission to multiple IAM users
B. AWS S3 will verify proper rights given by the owner of Account A, the bucket owner as well as by the IAM user B to the object
C. The bucket policy may not be created as S3 will give error due to conflict of Access Rights
D. It is not possible that the IAM user of one account accesses objects of the other IAM user

Answer: B

Explanation:
If a IAM user is trying to perform some action on an 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.html

NEW QUESTION 42
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/latest/UserGuide/VPC_NAT_Instance.htm|#EIP_Disable_Src DestCheck

NEW QUESTION 45
A user is thinking to use EBS PIOPS volume. Which of the below mentioned options is a right use case for the PIOPS EBS volume?

A. Analytics
B. System boot volume
C. Nlongo DB
D. Log processing

Answer: C

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. 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 business applications, database workloads, such as NoSQL DB, RDBMS, etc. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

NEW QUESTION 47
Which of the following is true of an instance profile when an IAM role is created using the console?

A. The instance profile uses a different name.
B. The console gives the instance profile the same name as the role it corresponds to.
C. The instance profile should be created manually by a user.
D. The console creates the role and instance profile as separate actions.

Answer: B

Explanation:
Amazon EC2 uses an instance profile as a container for an IAM role. When you create an IAM role using the console, the console creates an instance profile automatically and gives it the same name as the role it corresponds to. If you use the AWS CLI, API, or an AWS SDK to create a role, you create the role and instance profile as separate actions, and you might give them different names.
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2_instance-profiles.html

NEW QUESTION 48
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 50
AWS Direct Connect itself has NO specific resources for you to control access to. Therefore, there are no AWS Direct Connect Amazon Resource Names (ARNs) for you to use in an Identity and Access Nlanagement (IAM) policy. With that in mind, how is it possible to write a policy to control access to AWS Direct Connect actions?

A. You can leave the resource name field blank.
B. You can choose the name of the AWS Direct Connection as the resource.
C. You can use an asterisk (*) as the resource.
D. You can create a name for the resourc

Answer: C

Explanation:
AWS Direct Connect itself has no specific resources for you to control access to. Therefore, there are no AWS Direct Connect ARNs for you to use in an IAM policy. You use an asterisk (*) as the resource when writing a policy to control access to AWS Direct Connect actions.
Reference: http://docs.aws.amazon.com/directconnect/latest/UserGuide/using_iam.html

NEW QUESTION 51
With respect to AWS Lambda permissions model, at the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the role.

A. configuration
B. execution
C. delegation
D. dependency

Answer: B

Explanation:
Regardless of how your Lambda function is invoked, AWS Lambda always executes the function. At the time you create a Lambda function, you specify an IAM role that AWS Lambda can assume to execute your Lambda function on your behalf. This role is also referred to as the execution role.
Reference: http://docs.aws.amazon.com/Iambda/latest/dg/lambda-dg.pdf

NEW QUESTION 56
Within an IAM policy, can you add an IfExists condition at the end of a Null condition?

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

Answer: C

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

NEW QUESTION 60
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.html

NEW QUESTION 61
ExamKiIIer has created a multi-tenant Learning Management System (LMS). The application is hosted for five different tenants (clients) in the VPCs of the respective AWS accounts of the tenant. ExamKiIIer wants to setup a centralized server which can connect with the LMS of each tenant upgrade if required. ExamKiIIer also wants to ensure that one tenant VPC should not be able to connect to the other tenant VPC for security reasons. How can ExamKiIIer setup this scenario?

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

Answer:

A

Explanation:
A Virtual Private Cloud (VPC) is a virtual network dedicated to the user's AWS account. It enables the user to launch AWS resources into a virtual network that the user has defined. A VPC peering connection allows the user to route traffic between the peer VPCs using private IP addresses as if they are a part of the same network.
This is helpful when one VPC from the same or different AWS account wants to connect with resources of the other VPC. The organization wants to setup that one VPC can connect with all the other VPCs but all other VPCs cannot connect among each other. This can be achieved by configuring VPC peering where one VPC is peered with all the other VPCs, but the other VPCs are not peered to each other. The VPCs are in the same or a separate AWS account and should not have overlapping CIDR blocks.
Reference:
http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/peering-configurations-full-access.html# many-vpcs-full-acces

NEW QUESTION 65
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 67
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.htmI

NEW QUESTION 68
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 70
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 72
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.html

NEW QUESTION 75
An organization is setting up a web application with the JEE stack. The application uses the JBoss app server and |V|ySQL DB. The application has a logging module which logs all the actMties whenever a business function of the JEE application is called. The logging actMty takes some time due to the large size of the log file. If the application wants to setup a scalable infrastructure which of the below mentioned options will help achieve this setup?

A. Host the log files on EBS with PIOPS which will have higher I/O.
B. Host logging and the app server on separate sewers such that they are both in the same zone.
C. Host logging and the app server on the same instance so that the network latency will be shorter.
D. Create a separate module for logging and using SQS compartmentalize the module such that all calls to logging are asynchronous.

Answer: D

Explanation:
The organization can always launch multiple EC2 instances in the same region across multiple AZs for HA and DR. The AWS architecture practice recommends compartmentalizing the functionality such that
they can both run in parallel without affecting the performance of the main application. In this scenario logging takes a longer time due to the large size of the log file. Thus, it is recommended that the organization should separate them out and make separate modules and make asynchronous calls among them. This way the application can scale as per the requirement and the performance will not bear the impact of logging.
Reference: http://www.awsarchitectureblog.com/2014/03/aws-and-compartmentalization.html

NEW QUESTION 78
A user has set the IAM policy where it denies all requests if a request is not from IP 10.10.10.1/32. The other policy says allow all requests between 5 PM to 7 PM. What will happen when a user is requesting access from IP 55.109.10.12/32 at 6 PM?

A. It will deny access
B. It is not possible to set a policy based on the time or IP
C. IAM will throw an error for policy conflict
D. It will allow access

Answer: A

Explanation:
When a request is made, the AWS IAM policy decides whether a given request should be allowed or denied. The evaluation logic follows these rules:
By default, all requests are denied. (In general, requests made using the account credentials for resources in the account are always allowed.)
An explicit allow policy overrides this default.
An explicit deny policy overrides any allows.
In this case since there are explicit deny and explicit allow statements. Thus, the request will be denied since deny overrides allow.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/AccessPolicyLanguage_EvaluationLogic.html

NEW QUESTION 80
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.html#createvirtualinterface

NEW QUESTION 84
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 88
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 92
Identify a true statement about the statement ID (Sid) in IAM.

A. You cannot expose the Sid in the IAM API.
B. You cannot use a Sid value as a sub-ID for a policy document's ID for services provided by SQS and SNS.
C. You can expose the Sid in the IAM API.
D. You cannot assign a Sid value to each statement in a statement arra

Answer: A

Explanation:
The Sid(statement ID) is an optional identifier that you provide for the policy statement. You can assign a Sid a value to each statement in a statement array. In IAM, the Sid is not exposed in the IAM API. You can't retrieve a particular statement based on this ID.
Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements.html#Sid

NEW QUESTION 95
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, E|astiCache 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/AmazonE|astiCache/latest/UserGuide/AmazonVPC.EC.html

NEW QUESTION 98
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 .large instances.
E. Route traffic to EC2 m1 .large 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 102
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 107
You are tasked with moving a legacy application from a virtual machine running Inside your datacenter to an Amazon VPC Unfortunately this app requires access to a number of on-premises services and no one who configured the app still works for your company. Even worse there's no documentation for it.
What will allow the application running inside the VPC to reach back and access its internal dependencies
without being reconfigured? (Choose 3 answers)

A. An AWS Direct Connect link between the VPC and the network housing the internal services.
B. An Internet Gateway to allow a VPN connection.
C. An Elastic IP address on the VPC instance
D. An IP address space that does not conflict with the one on-premises
E. Entries in Amazon Route 53 that allow the Instance to resolve its dependencies' IP addresses
F. A VM Import of the current virtual machine

Answer: ADF

NEW QUESTION 108
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 111
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 115
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 116
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 118
You are designing an intrusion detection prevention (IDS/IPS) solution for a customer web application in a single VPC. You are considering the options for implementing IOS IPS protection for traffic coming from the Internet.
Which of the following options would you consider? (Choose 2 answers)

A. Implement IDS/IPS agents on each Instance running In VPC
B. Configure an instance in each subnet to switch its network interface card to promiscuous mode and analyze network traffic.
C. Implement Elastic Load Balancing with SSL listeners In front of the web applications
D. Implement a reverse proxy layer in front of web servers and configure IDS/IPS agents on each reverse proxy server.

Answer: BD

NEW QUESTION 122
Your company previously configured a heavily used, dynamically routed VPN connection between your on-premises data center and AWS. You recently provisioned a DirectConnect connection and would like to start using the new connection. After configuring DirectConnect settings in the AWS Console, which of the following options win provide the most seamless transition for your users?

A. Delete your existing VPN connection to avoid routing loops configure your DirectConnect router with the appropriate settings and verity network traffic is leveraging DirectConnect.
B. Configure your DirectConnect router with a higher BGP priority man your VPN router, verify network traffic is leveraging Directconnect and then delete your existing VPN connection.
C. Update your VPC route tables to point to the DirectConnect connection configure your DirectConnect router with the appropriate settings verify network traffic is leveraging DirectConnect and then delete the VPN connection.
D. Configure your DirectConnect router, update your VPC route tables to point to the DirectConnect connection, configure your VPN connection with a higher BGP priorit
E. And verify network traffic is leveraging the DirectConnect connection.

Answer: D

NEW QUESTION 124
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 128
You are migrating a legacy client-server application to AWS. The application responds to a specific DNS domain (e.g. www.exampI|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 MuIti-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 130
You are designing a data leak prevention solution for your VPC environment. You want your VPC Instances to be able to access software depots and distributions on the Internet for product updates. The depots and distributions are accessible via third party CDNs by their URLs. You want to explicitly deny any other outbound connections from your VPC instances to hosts on the internet.
Which of the following options would you consider?

A. Configure a web proxy server in your VPC and enforce URL-based rules for outbound access Remove default routes.
B. Implement security groups and configure outbound rules to only permit traffic to software depots.
C. Move all your instances into private VPC subnets remove default routes from all routing tables and add specific routes to the software depots and distributions only.
D. Implement network access control lists to all specific destinations, with an Implicit deny as a rul

Answer: A

NEW QUESTION 132
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 136
Your company has recently extended its datacenter into a VPC on AWS to add burst computing capacity as needed Members of your Network Operations Center need to be able to go to the AWS Management Console and administer Amazon EC2 instances as necessary You don't want to create new IAM users for each NOC member and make those users sign in again to the AWS Management Console Which option below will meet the needs for your NOC members?

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

Answer: D

NEW QUESTION 140
You are developing a new mobile application and are considering storing user preferences in AWS.2w This would provide a more uniform cross-device experience to users using multiple mobile devices to access the application. The preference data for each user is estimated to be 50KB in size Additionally 5 million customers are expected to use the application on a regular basis. The solution needs to be
cost-effective, highly available, scalable and secure, how would you design a solution to meet the above requirements?

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

Answer: B

NEW QUESTION 142
In AWS, which security aspects are the customer's responsibility? Choose 4 answers

A. Security Group and ACL (Access Control List) settings
B. Decommissioning storage devices
C. Patch management on the EC2 instance's operating system
D. Life-cycle management of IAM credentials
E. Controlling physical access to compute resources

F. Encryption of EBS (Elastic Block Storage) volumes

Answer: ACDF

NEW QUESTION 144

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

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

Answer: A

NEW QUESTION 147

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 152

Your company has an on-premises multi-tier PHP web application, which recently experienced downtime due to a large burst In web traffic due to a company announcement Over the coming days, you are expecting similar announcements to drive similar unpredictable bursts, and are looking to find ways to quickly improve your infrastructures ability to handle unexpected increases in traffic.
The application currently consists of 2 tiers a web tier which consists of a load balancer and several Linux Apache web servers as well as a database tier which hosts a Linux server hosting a MySQL database. Which scenario below will provide full site functionality, while helping to improve the ability of your application in the short timeframe required?

A. Failover environment: Create an S3 bucket and configure it for website hostin
B. Migrate your DNS to Route53 using zone file import, and leverage Route53 DNS failover to failover to the S3 hosted website.
C. Hybrid environment: Create an AMI, which can be used to launch web sewers in EC2. Create an Auto Scaling group, which uses the AMI to scale the web tier based on incoming traffi
D. Leverage Elastic Load Balancing to balance traffic between on-premises web servers and those hosted In AWS.
E. Offload traffic from on-premises environment: Setup a CloudFront distribution, and configure CloudFront to cache objects from a custom origi
F. Choose to customize your object cache behavior, and select a TTL that objects should exist in cache.
G. Migrate to AWS: Use VM Import/Export to quickly convert an on-premises web server to an AM
H. Create an Auto Scaling group, which uses the imported AMI to scale the web tier based on incoming traffi
I. Create an RDS read replica and setup replication between the RDS instance and on-premises MySQL server to migrate the database.

Answer: C

NEW QUESTION 153

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 156

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 158
You control access to S3 buckets and objects with:

A. Identity and Access Management (IAM) Policies.
B. Access Control Lists (ACLs).
C. Bucket Policies.
D. All of the above

Answer: D

NEW QUESTION 162
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 164
......

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