

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 1 – Question 50**

QUESTION 1

Amazon SWF is designed to help users_____.

- A. Design graphical user interface interactions.
- B. Manage user identification and authorization.
- C. Store Web content.
- D. Coordinate synchronous and asynchronous tasks which are distributed and fault tolerant.

Answer: D

QUESTION 2

Can I control if and when MySQL based RDS Instance is upgraded to new supported versions?

- A. No
- B. Only in VPC
- C. Yes

Answer: C

QUESTION 3

What does Amazon Elastic Beanstalk provide?

- A. A scalable storage appliance on top of Amazon Web Services.
- B. An application container on top of Amazon Web Services.
- C. A service by this name doesn't exist.
- D. A scalable cluster of EC2 instances.

Answer: B

QUESTION 4

True or False: When using IAM to control access to your RDS resources, the key names that can be used are case sensitive. For example, aws:CurrentTime is NOT equivalent to AWS:currenttime.

- A. TRUE
- B. FALSE

Answer: B

Explanation:

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

Key names are case-insensitive. For example, aws:CurrentTime is equivalent to AWS:currenttime.

QUESTION 5

What will be the status of the snapshot until the snapshot is complete.

- A. running
- B. working
- C. progressing
- D. pending

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

QUESTION 6

Can we attach an EBS volume to more than one EC2 instance at the same time?

- A. No
- B. Yes.
- C. Only EC2-optimized EBS volumes.
- D. Only in read mode.

Answer: A

QUESTION 7

True or False: Automated backups are enabled by default for a new DB Instance.

- A. TRUE
- B. FALSE

Answer: A

QUESTION 8

What does the AWS Storage Gateway provide?

- A. It allows to integrate on-premises IT environments with Cloud Storage.
- B. A direct encrypted connection to Amazon S3.
- C. It's a backup solution that provides an on-premises Cloud storage.
- D. It provides an encrypted SSL endpoint for backups in the Cloud.

Answer: A

QUESTION 9

Amazon RDS automated backups and DB Snapshots are currently supported for only the ____ storage engine.

- A. InnoDB
- B. MyISAM

Answer: A

QUESTION 10

How many relational database engines does RDS currently support?

- A. Three: MySQL, Oracle and Microsoft SQL Server.
- B. Just two: MySQL and Oracle.
- C. Six: Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB.
- D. Just one: MySQL.

Answer: A

Explanation:

Amazon RDS provides you six familiar database engines to choose from, including Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB.

https://aws.amazon.com/rds/?nc1=h_ls

QUESTION 11

Fill in the blanks: The base URI for all requests for instance metadata is _____.

- A. <http://254.169.169.254/latest/>
- B. <http://169.169.254.254/latest/>
- C. <http://127.0.0.1/latest/>
- D. <http://169.254.169.254/latest/>

Answer: D

Explanation:

<http://aws.amazon.com/search?searchQuery=metadata&searchPath=all&x=0&y=0>

QUESTION 12

While creating the snapshots using the command line tools, which command should I be using?

- A. `ec2-deploy-snapshot`
- B. `ec2-fresh-snapshot`
- C. `ec2-create-snapshot`
- D. `ec2-new-snapshot`

Answer: C

Explanation:

<http://docs.aws.amazon.com/cli/latest/reference/ec2/create-snapshot.html>

QUESTION 13

Typically, you want your application to check whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an _____ node in the response from the Amazon RDS API.

- A. Incorrect
- B. Error
- C. FALSE

Answer: B

Explanation:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/APITroubleshooting.html>

QUESTION 14

What are the two types of licensing options available for using Amazon RDS for Oracle?

- A. BYOL and Enterprise License
- B. BYOL and License Included
- C. Enterprise License and License Included
- D. Role based License and License Included

Answer: B

Explanation:

<https://aws.amazon.com/rds/oracle/>

QUESTION 15

What does a "Domain" refer to in Amazon SWF?

- A. A security group in which only tasks inside can communicate with each other
- B. A special type of worker
- C. A collection of related Workflows
- D. The DNS record for the Amazon SWF service

Answer: C

Explanation:

Domains provide a way of scoping Amazon SWF resources within your AWS account. All the components of a workflow, such as the workflow type and activity types, must be specified to be in a domain. It is possible to have more than one workflow in a domain; however, workflows in different domains cannot interact with each other.

<http://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dev-domain.html>

QUESTION 16

EBS Snapshots occur _____.

- A. Asynchronously
- B. Synchronously
- C. Weekly

Answer: A

Explanation:

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

QUESTION 17

Disabling automated backups _____ disable the point-in-time recovery.

- A. if configured to can
- B. will never
- C. will

Answer: C

QUESTION 18

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

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

Answer: C

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/raid-config.html>

raid 0 and 1 are the common types. Raid 5 and 6 are not recommended because of the extended stripe. If you encounter this question on the exam I suspect the answer options will be different.

Raid 1 Disadvantage

Does not provide a write performance improvement; requires more Amazon EC2 to Amazon EBS bandwidth than non-RAID configurations because the data is written to multiple volumes simultaneously.

Raid 0 Disadvantage

Performance of the stripe is limited to the worst performing volume in the set. Loss of a single volume results in a complete data loss for the array.

Raid 5 and 6 notes

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. Depending on the configuration of your RAID array, these RAID modes provide 20-30% fewer usable IOPS than a RAID 0 configuration. Increased cost is a factor with these RAID modes as well; when using identical volume sizes and speeds, a 2-volume RAID 0 array can outperform a 4-volume RAID 6 array that costs twice as much.

QUESTION 19

Is creating a Read Replica of another Read Replica supported?

- A. Only in certain regions
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS types
- D. No

Answer: B

Explanation:

<https://aws.amazon.com/rds/faqs/>

Q: Can I create a Read Replica of another Read Replica?

Amazon RDS for MySQL: You can create a second-tier Read Replica from an existing first-tier Read Replica. By creating a second-tier Read Replica, you may be able to move some of the replication load from the master database instance to a first-tier Read Replica. Please note that a second-tier Read Replica may lag further behind the master because of additional replication latency introduced as transactions are replicated from the master to the first tier replica and then to the second-tier replica.

Amazon RDS for PostgreSQL: Read Replicas of Read Replicas are not currently supported.

QUESTION 20

Can Amazon S3 uploads resume on failure or do they need to restart?

- A. Restart from beginning
- B. You can resume them, if you flag the "resume on failure" option before uploading.
- C. Resume on failure
- D. Depends on the file size

Answer: C

QUESTION 21

Which of the following cannot be used in Amazon EC2 to control who has access to specific Amazon EC2 instances?

- A. Security Groups
- B. IAM System
- C. SSH keys
- D. Windows passwords

Answer: B

Explanation:

<http://blogs.aws.amazon.com/security/post/Tx29HCT3ABL7LP3/Resource-level-Permissions-for-EC2-Controlling-Management-Access-on-Specific-Ins>

QUESTION 22

Fill in the blanks: _____ let you categorize your EC2 resources in different ways, for example, by purpose, owner, or environment.

- A. wildcards
- B. pointers
- C. Tags
- D. special filters

Answer: C

QUESTION 23

How can I change the security group membership for interfaces owned by other AWS, such as Elastic Load Balancing?

- A. By using the service specific console or API/CLI commands
- B. None of these
- C. Using Amazon EC2 API/CLI
- D. using all these methods

Answer: A

Explanation:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-security-groups.html>

QUESTION 24

What is the maximum write throughput I can provision for a single Dynamic DB table?

- A. 1,000 write capacity units
- B. 100,000 write capacity units
- C. Dynamic DB is designed to scale without limits, but if you go beyond 10,000 you have to contact AWS first.
- D. 10,000 write capacity units

Answer: C

Explanation:

<https://aws.amazon.com/dynamodb/faqs/>

QUESTION 25

What does the following command do with respect to the Amazon EC2 security groups? `ec2-revoke RevokeSecurityGroupIngress`.

- A. Removes one or more security groups from a rule.
- B. Removes one or more security groups from an Amazon EC2 instance.
- C. Removes one or more rules from a security group.
- D. Removes a security group from our account.

Answer: C

Explanation:

Removes one or more ingress rules from a security group. The values that you specify in the revoke request (for example, ports) must match the existing rule's values for the rule to be removed.

<http://docs.aws.amazon.com/cli/latest/reference/ec2/revoke-security-group-ingress.html>

QUESTION 26

Can a 'user' be associated with multiple AWS accounts?

- A. No
- B. Yes

Answer: A

QUESTION 27

True or False: Manually created DB Snapshots are deleted after the DB Instance is deleted.

- A. TRUE
- B. FALSE

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_DeleteInstance.html

If you choose not to create a final DB snapshot, you will not be able to later restore the DB instance to its final state. When you delete a DB instance, all automated backups are deleted and cannot be recovered. Manual DB snapshots of the instance are not deleted.

QUESTION 28

What happens to the data on an instance if the instance reboots (intentionally or unintentionally)?

- A. Data will be lost
- B. Data persists
- C. Data may persist however cannot be sure

Answer: B

Explanation:

Instance Store Lifetime

You can specify instance store volumes for an instance only when you launch it. The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data in the instance store is lost under the following circumstances:

The underlying disk drive fails

The instance stops

The instance terminates

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html>

QUESTION 29

How many types of block devices does Amazon EC2 support?

- A. 2

- B. 3
- C. 4
- D. 1

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

Amazon EC2 supports two types of block devices:

Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)

EBS volumes (remote storage devices)

A block device mapping defines the block devices (instance store volumes and EBS volumes) to attach to an instance.

QUESTION 30

Provisioned IOPS Costs: you are charged for the IOPS and storage whether or not you use them in a given month.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

EBS Provisioned IOPS SSD (io1) Volumes

Volume storage for EBS Provisioned IOPS SSD (io1) volumes is charged by the amount you provision in GB per month, until you release the storage. With Provisioned IOPS SSD (io1) volumes, you are also charged by the amount you provision in IOPS (input/output operations per second) multiplied by the percentage of days you provision for the month. For example, if you provision a volume with 1000 IOPS, and keep this volume for 15 days in a 30 day month, then in a Region that charges \$0.10 per provisioned IOPS-month, you would be charged \$50 for the IOPS that you provision ($\$0.10 \text{ per provisioned IOPS-month} * 1000 \text{ IOPS provisioned} * 15 \text{ days}/30$). You will be charged for the IOPS provisioned on a volume even when the volume is detached from an instance.

<https://aws.amazon.com/ebs/pricing/>

QUESTION 31

IAM provides several policy templates you can use to automatically assign permissions to the groups you create. The ____ policy template gives the Admins group permission to access all account resources, except your AWS account information

- A. Read Only Access
- B. Power User Access
- C. AWS Cloud Formation Read Only Access
- D. Administrator Access

Answer: B

Explanation:

AWS managed policies are designed to provide permissions for many common use cases. For example, there are AWS managed policies that define typical permissions for administrators (all access), for power users (all access except IAM), and for other various levels of access to AWS services. AWS managed policies make it easier for you to assign appropriate permissions to users, groups, and roles than if you had to write the policies yourself.

http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_managed-vs-inline.html

QUESTION 32

While performing the volume status checks, if the status is insufficient-data, what does it mean?

- A. the checks may still be in progress on the volume
- B. the check has passed
- C. the check has failed

Answer: A

Explanation:

If the status is insufficient-data, the checks may still be in progress on the volume. You can view the results of volume status checks to identify any impaired volumes and take any necessary actions.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html#monitoring-volume-checks>

QUESTION 33

IAM's Policy Evaluation Logic always starts with a default ____ for every request, except for those that use the AWS account's root security credentials by default.

- A. Permit
- B. Deny
- C. Cancel

Answer: B

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html

QUESTION 34

By default, when an EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the ____ Service to set the drive letters of the EBS volumes per your specifications.

- A. EBSConfig Service
- B. AMIConfig Service
- C. Ec2Config Service
- D. Ec2-AMIConfig Service

Answer: C

Explanation:

Ec2Config Service is like sysprep and used specifically for windows instances.. You can change parameters in OS before launching.

QUESTION 35

For each DB Instance class, what is the maximum size of associated storage capacity?

- A. 5GB
- B. 1TB
- C. 6TB
- D. 500GB

Answer: C

Explanation:

"You can now create MySQL, PostgreSQL, and Oracle RDS database instances with up to 6TB of storage and SQL Server RDS database instances with up to 4TB of storage when using the Provisioned IOPS and General Purpose (SSD) storage types. Existing MySQL, PostgreSQL, and Oracle RDS database instances can be scaled to these new database storage limits without any downtime."

QUESTION 36

SQL Server ____ store logins and passwords in the master database.

- A. can be configured to but by default does not
- B. doesn't
- C. does

Answer: C

Explanation:

There are two authentications

Windows authentication

The credentials for which are not stored in SQL Server database and managed by windows/AD. There would be entry for windows authenticated logins in master database with respective SID but password would be with Active directory.

SQL Server authentication.

For 2nd we have password stored in hash format you can see it from sys.sql_logins. The information about SQL server logins are stored in master database and each login has SID receptive to it. Only SA login has same SID no matter what server it is. That is why when you move database by backup restore mechanism users are moved not logins and you finally have to create logins(if already not there) and map it to users. This is generally called as troubleshooting orphaned users

QUESTION 37

What is Oracle SQL Developer?

- A. An AWS developer who is an expert in Amazon RDS using both the Oracle and SQL Server DB engines
- B. A graphical Java tool distributed without cost by Oracle
- C. It is a variant of the SQL Server Management Studio designed by Microsoft to support Oracle DBMS functionalities
- D. A different DBMS released by Microsoft free of cost

Answer: B

Explanation:

<http://www.oracle.com/technetwork/developer-tools/sql-developer/what-is-sqldev-093866.html>

QUESTION 38

Does Amazon RDS allow direct host access via Telnet, Secure Shell (SSH), or Windows Remote Desktop Connection?

- A. Yes
- B. No
- C. Depends on if it is in VPC or not

Answer: B

QUESTION 39

To view information about an Amazon EBS volume, open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>, click ____ in the Navigation pane.

- A. EBS
- B. Describe
- C. Details
- D. Volumes

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-describing-volumes.html>

QUESTION 40

Using Amazon IAM, can I give permission based on organizational groups?

- A. Yes but only in certain cases
- B. No
- C. Yes always

Answer: C

Explanation:

An IAM group is a collection of IAM users. You can use groups to specify permissions for a collection of users, which can make those permissions easier to manage for those users.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/id.html>

QUESTION 41

While creating the snapshots using the API, which Action should I be using?

- A. MakeSnapShot
- B. FreshSnapshot
- C. DeploySnapshot
- D. CreateSnapshot

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd-CreateSnapshot.html>

QUESTION 42

What is an isolated database environment running in the cloud (Amazon RDS) called?

- A. DB Instance

- B. DB Server
- C. DB Unit
- D. DB Volume

Answer: A

QUESTION 43

While signing in REST/ Query requests, for additional security, you should transmit your requests using Secure Sockets Layer (SSL) by using _____.

- A. HTTP
- B. Internet Protocol Security(IPsec)
- C. TLS (Transport Layer Security)
- D. HTTPS

Answer: D

QUESTION 44

What happens to the I/O operations while you take a database snapshot?

- A. I/O operations to the database are suspended for a few minutes while the backup is in progress.
- B. I/O operations to the database are sent to a Replica (if available) for a few minutes while the backup is in progress.
- C. I/O operations will be functioning normally.
- D. I/O operations to the database are suspended for an hour while the backup is in progress.

Answer: A

Explanation:

Creating this DB snapshot on a Single-AZ DB instance results in a brief I/O suspension that typically lasting no more than a few minutes. Multi-AZ DB instances are not affected by this I/O suspension since the backup is taken on the standby.

QUESTION 45

Read Replicas require a transactional storage engine and are only supported for the _____ storage engine.

- A. OracleISAM
- B. MSSQLDB
- C. InnoDB
- D. MyISAM

Answer: C

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 46

When running my DB Instance as a Multi-AZ deployment, can I use the standby for read or write operations?

- A. Yes
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS instances
- D. No

Answer: D

QUESTION 47

When should I choose Provisioned IOPS over Standard RDS storage?

- A. If you have batch-oriented workloads
- B. If you use production online transaction processing (OLTP) workloads
- C. If you have workloads that are not sensitive to consistent performance

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

Amazon RDS provisions that IOPS rate and storage for the lifetime of the DB instance or until you change it. Provisioned IOPS storage is optimized for I/O intensive, online transaction processing (OLTP) workloads that have consistent performance requirements. Provisioned IOPS helps performance tuning.

QUESTION 48

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send ____ minute metrics to Amazon CloudWatch.

- A. 3
- B. 1
- C. 5
- D. 2

Answer: B

QUESTION 49

What is the minimum charge for the data transferred between Amazon RDS and Amazon EC2 Instances in the same Availability Zone?

- A. USD 0.10 per GB
- B. No charge. It is free.
- C. USD 0.02 per GB
- D. USD 0.01 per GB

Answer: B

QUESTION 50

Are Reserved Instances available for Multi-AZ Deployments?

- A. Only for Cluster Compute instances
- B. Yes for all instance types
- C. Only for M3 instance types
- D. No

Answer: B

Explanation:

<https://aws.amazon.com/rds/faqs/>

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 1 – Question 50**

QUESTION 1

Amazon SWF is designed to help users_____.

- A. Design graphical user interface interactions.
- B. Manage user identification and authorization.
- C. Store Web content.
- D. Coordinate synchronous and asynchronous tasks which are distributed and fault tolerant.

Answer: D

QUESTION 2

Can I control if and when MySQL based RDS Instance is upgraded to new supported versions?

- A. No
- B. Only in VPC
- C. Yes

Answer: C

QUESTION 3

What does Amazon Elastic Beanstalk provide?

- A. A scalable storage appliance on top of Amazon Web Services.
- B. An application container on top of Amazon Web Services.
- C. A service by this name doesn't exist.
- D. A scalable cluster of EC2 instances.

Answer: B

QUESTION 4

True or False: When using IAM to control access to your RDS resources, the key names that can be used are case sensitive. For example, aws:CurrentTime is NOT equivalent to AWS:currenttime.

- A. TRUE
- B. FALSE

Answer: B

Explanation:

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

Key names are case-insensitive. For example, aws:CurrentTime is equivalent to AWS:currenttime.

QUESTION 5

What will be the status of the snapshot until the snapshot is complete.

- A. running
- B. working
- C. progressing
- D. pending

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

QUESTION 6

Can we attach an EBS volume to more than one EC2 instance at the same time?

- A. No
- B. Yes.
- C. Only EC2-optimized EBS volumes.
- D. Only in read mode.

Answer: A

QUESTION 7

True or False: Automated backups are enabled by default for a new DB Instance.

- A. TRUE
- B. FALSE

Answer: A

QUESTION 8

What does the AWS Storage Gateway provide?

- A. It allows to integrate on-premises IT environments with Cloud Storage.
- B. A direct encrypted connection to Amazon S3.
- C. It's a backup solution that provides an on-premises Cloud storage.
- D. It provides an encrypted SSL endpoint for backups in the Cloud.

Answer: A

QUESTION 9

Amazon RDS automated backups and DB Snapshots are currently supported for only the ____ storage engine.

- A. InnoDB
- B. MyISAM

Answer: A

QUESTION 10

How many relational database engines does RDS currently support?

- A. Three: MySQL, Oracle and Microsoft SQL Server.
- B. Just two: MySQL and Oracle.
- C. Six: Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB.
- D. Just one: MySQL.

Answer: A

Explanation:

Amazon RDS provides you six familiar database engines to choose from, including Amazon Aurora, Oracle, Microsoft SQL Server, PostgreSQL, MySQL and MariaDB.

https://aws.amazon.com/rds/?nc1=h_ls

QUESTION 11

Fill in the blanks: The base URI for all requests for instance metadata is _____.

- A. <http://254.169.169.254/latest/>
- B. <http://169.169.254.254/latest/>
- C. <http://127.0.0.1/latest/>
- D. <http://169.254.169.254/latest/>

Answer: D

Explanation:

<http://aws.amazon.com/search?searchQuery=metadata&searchPath=all&x=0&y=0>

QUESTION 12

While creating the snapshots using the command line tools, which command should I be using?

- A. `ec2-deploy-snapshot`
- B. `ec2-fresh-snapshot`
- C. `ec2-create-snapshot`
- D. `ec2-new-snapshot`

Answer: C

Explanation:

<http://docs.aws.amazon.com/cli/latest/reference/ec2/create-snapshot.html>

QUESTION 13

Typically, you want your application to check whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an _____ node in the response from the Amazon RDS API.

- A. Incorrect
- B. Error
- C. FALSE

Answer: B

Explanation:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/APITroubleshooting.html>

QUESTION 14

What are the two types of licensing options available for using Amazon RDS for Oracle?

- A. BYOL and Enterprise License
- B. BYOL and License Included
- C. Enterprise License and License Included
- D. Role based License and License Included

Answer: B

Explanation:

<https://aws.amazon.com/rds/oracle/>

QUESTION 15

What does a "Domain" refer to in Amazon SWF?

- A. A security group in which only tasks inside can communicate with each other
- B. A special type of worker
- C. A collection of related Workflows
- D. The DNS record for the Amazon SWF service

Answer: C

Explanation:

Domains provide a way of scoping Amazon SWF resources within your AWS account. All the components of a workflow, such as the workflow type and activity types, must be specified to be in a domain. It is possible to have more than one workflow in a domain; however, workflows in different domains cannot interact with each other.

<http://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dev-domain.html>

QUESTION 16

EBS Snapshots occur _____.

- A. Asynchronously
- B. Synchronously
- C. Weekly

Answer: A

Explanation:

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

QUESTION 17

Disabling automated backups _____ disable the point-in-time recovery.

- A. if configured to can
- B. will never
- C. will

Answer: C

QUESTION 18

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

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

Answer: C

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/raid-config.html>

raid 0 and 1 are the common types. Raid 5 and 6 are not recommended because of the extended stripe. If you encounter this question on the exam I suspect the answer options will be different.

Raid 1 Disadvantage

Does not provide a write performance improvement; requires more Amazon EC2 to Amazon EBS bandwidth than non-RAID configurations because the data is written to multiple volumes simultaneously.

Raid 0 Disadvantage

Performance of the stripe is limited to the worst performing volume in the set. Loss of a single volume results in a complete data loss for the array.

Raid 5 and 6 notes

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. Depending on the configuration of your RAID array, these RAID modes provide 20-30% fewer usable IOPS than a RAID 0 configuration. Increased cost is a factor with these RAID modes as well; when using identical volume sizes and speeds, a 2-volume RAID 0 array can outperform a 4-volume RAID 6 array that costs twice as much.

QUESTION 19

Is creating a Read Replica of another Read Replica supported?

- A. Only in certain regions
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS types
- D. No

Answer: B

Explanation:

<https://aws.amazon.com/rds/faqs/>

Q: Can I create a Read Replica of another Read Replica?

Amazon RDS for MySQL: You can create a second-tier Read Replica from an existing first-tier Read Replica. By creating a second-tier Read Replica, you may be able to move some of the replication load from the master database instance to a first-tier Read Replica. Please note that a second-tier Read Replica may lag further behind the master because of additional replication latency introduced as transactions are replicated from the master to the first tier replica and then to the second-tier replica.

Amazon RDS for PostgreSQL: Read Replicas of Read Replicas are not currently supported.

QUESTION 20

Can Amazon S3 uploads resume on failure or do they need to restart?

- A. Restart from beginning
- B. You can resume them, if you flag the "resume on failure" option before uploading.
- C. Resume on failure
- D. Depends on the file size

Answer: C

QUESTION 21

Which of the following cannot be used in Amazon EC2 to control who has access to specific Amazon EC2 instances?

- A. Security Groups
- B. IAM System
- C. SSH keys
- D. Windows passwords

Answer: B

Explanation:

<http://blogs.aws.amazon.com/security/post/Tx29HCT3ABL7LP3/Resource-level-Permissions-for-EC2-Controlling-Management-Access-on-Specific-Ins>

QUESTION 22

Fill in the blanks: _____ let you categorize your EC2 resources in different ways, for example, by purpose, owner, or environment.

- A. wildcards
- B. pointers
- C. Tags
- D. special filters

Answer: C

QUESTION 23

How can I change the security group membership for interfaces owned by other AWS, such as Elastic Load Balancing?

- A. By using the service specific console or API/CLI commands
- B. None of these
- C. Using Amazon EC2 API/CLI
- D. using all these methods

Answer: A

Explanation:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-security-groups.html>

QUESTION 24

What is the maximum write throughput I can provision for a single Dynamic DB table?

- A. 1,000 write capacity units
- B. 100,000 write capacity units
- C. Dynamic DB is designed to scale without limits, but if you go beyond 10,000 you have to contact AWS first.
- D. 10,000 write capacity units

Answer: C

Explanation:

<https://aws.amazon.com/dynamodb/faqs/>

QUESTION 25

What does the following command do with respect to the Amazon EC2 security groups? `ec2-revoke RevokeSecurityGroupIngress`.

- A. Removes one or more security groups from a rule.
- B. Removes one or more security groups from an Amazon EC2 instance.
- C. Removes one or more rules from a security group.
- D. Removes a security group from our account.

Answer: C

Explanation:

Removes one or more ingress rules from a security group. The values that you specify in the revoke request (for example, ports) must match the existing rule's values for the rule to be removed.

<http://docs.aws.amazon.com/cli/latest/reference/ec2/revoke-security-group-ingress.html>

QUESTION 26

Can a 'user' be associated with multiple AWS accounts?

- A. No
- B. Yes

Answer: A

QUESTION 27

True or False: Manually created DB Snapshots are deleted after the DB Instance is deleted.

- A. TRUE
- B. FALSE

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_DeleteInstance.html

If you choose not to create a final DB snapshot, you will not be able to later restore the DB instance to its final state. When you delete a DB instance, all automated backups are deleted and cannot be recovered. Manual DB snapshots of the instance are not deleted.

QUESTION 28

What happens to the data on an instance if the instance reboots (intentionally or unintentionally)?

- A. Data will be lost
- B. Data persists
- C. Data may persist however cannot be sure

Answer: B

Explanation:

Instance Store Lifetime

You can specify instance store volumes for an instance only when you launch it. The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data in the instance store is lost under the following circumstances:

The underlying disk drive fails

The instance stops

The instance terminates

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html>

QUESTION 29

How many types of block devices does Amazon EC2 support?

- A. 2

- B. 3
- C. 4
- D. 1

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/block-device-mapping-concepts.html>

Amazon EC2 supports two types of block devices:

Instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)

EBS volumes (remote storage devices)

A block device mapping defines the block devices (instance store volumes and EBS volumes) to attach to an instance.

QUESTION 30

Provisioned IOPS Costs: you are charged for the IOPS and storage whether or not you use them in a given month.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

EBS Provisioned IOPS SSD (io1) Volumes

Volume storage for EBS Provisioned IOPS SSD (io1) volumes is charged by the amount you provision in GB per month, until you release the storage. With Provisioned IOPS SSD (io1) volumes, you are also charged by the amount you provision in IOPS (input/output operations per second) multiplied by the percentage of days you provision for the month. For example, if you provision a volume with 1000 IOPS, and keep this volume for 15 days in a 30 day month, then in a Region that charges \$0.10 per provisioned IOPS-month, you would be charged \$50 for the IOPS that you provision ($\$0.10$ per provisioned IOPS-month * 1000 IOPS provisioned * 15 days/30). You will be charged for the IOPS provisioned on a volume even when the volume is detached from an instance.

<https://aws.amazon.com/ebs/pricing/>

QUESTION 31

IAM provides several policy templates you can use to automatically assign permissions to the groups you create. The ____ policy template gives the Admins group permission to access all account resources, except your AWS account information

- A. Read Only Access
- B. Power User Access
- C. AWS Cloud Formation Read Only Access
- D. Administrator Access

Answer: B

Explanation:

AWS managed policies are designed to provide permissions for many common use cases. For example, there are AWS managed policies that define typical permissions for administrators (all access), for power users (all access except IAM), and for other various levels of access to AWS services. AWS managed policies make it easier for you to assign appropriate permissions to users, groups, and roles than if you had to write the policies yourself.

http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_managed-vs-inline.html

QUESTION 32

While performing the volume status checks, if the status is insufficient-data, what does it mean?

- A. the checks may still be in progress on the volume
- B. the check has passed
- C. the check has failed

Answer: A

Explanation:

If the status is insufficient-data, the checks may still be in progress on the volume. You can view the results of volume status checks to identify any impaired volumes and take any necessary actions.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html#monitoring-volume-checks>

QUESTION 33

IAM's Policy Evaluation Logic always starts with a default ____ for every request, except for those that use the AWS account's root security credentials by default.

- A. Permit
- B. Deny
- C. Cancel

Answer: B

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html

QUESTION 34

By default, when an EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the ____ Service to set the drive letters of the EBS volumes per your specifications.

- A. EBSConfig Service
- B. AMIConfig Service
- C. Ec2Config Service
- D. Ec2-AMIConfig Service

Answer: C

Explanation:

Ec2Config Service is like sysprep and used specifically for windows instances.. You can change parameters in OS before launching.

QUESTION 35

For each DB Instance class, what is the maximum size of associated storage capacity?

- A. 5GB
- B. 1TB
- C. 6TB
- D. 500GB

Answer: C

Explanation:

"You can now create MySQL, PostgreSQL, and Oracle RDS database instances with up to 6TB of storage and SQL Server RDS database instances with up to 4TB of storage when using the Provisioned IOPS and General Purpose (SSD) storage types. Existing MySQL, PostgreSQL, and Oracle RDS database instances can be scaled to these new database storage limits without any downtime."

QUESTION 36

SQL Server ____ store logins and passwords in the master database.

- A. can be configured to but by default does not
- B. doesn't
- C. does

Answer: C

Explanation:

There are two authentications

Windows authentication

The credentials for which are not stored in SQL Server database and managed by windows/AD. There would be entry for windows authenticated logins in master database with respective SID but password would be with Active directory.

SQL Server authentication.

For 2nd we have password stored in hash format you can see it from sys.sql_logins. The information about SQL server logins are stored in master database and each login has SID receptive to it. Only SA login has same SID no matter what server it is. That is why when you move database by backup restore mechanism users are moved not logins and you finally have to create logins(if already not there) and map it to users. This is generally called as troubleshooting orphaned users

QUESTION 37

What is Oracle SQL Developer?

- A. An AWS developer who is an expert in Amazon RDS using both the Oracle and SQL Server DB engines
- B. A graphical Java tool distributed without cost by Oracle
- C. It is a variant of the SQL Server Management Studio designed by Microsoft to support Oracle DBMS functionalities
- D. A different DBMS released by Microsoft free of cost

Answer: B

Explanation:

<http://www.oracle.com/technetwork/developer-tools/sql-developer/what-is-sqldev-093866.html>

QUESTION 38

Does Amazon RDS allow direct host access via Telnet, Secure Shell (SSH), or Windows Remote Desktop Connection?

- A. Yes
- B. No
- C. Depends on if it is in VPC or not

Answer: B

QUESTION 39

To view information about an Amazon EBS volume, open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>, click ____ in the Navigation pane.

- A. EBS
- B. Describe
- C. Details
- D. Volumes

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-describing-volumes.html>

QUESTION 40

Using Amazon IAM, can I give permission based on organizational groups?

- A. Yes but only in certain cases
- B. No
- C. Yes always

Answer: C

Explanation:

An IAM group is a collection of IAM users. You can use groups to specify permissions for a collection of users, which can make those permissions easier to manage for those users.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/id.html>

QUESTION 41

While creating the snapshots using the API, which Action should I be using?

- A. MakeSnapShot
- B. FreshSnapshot
- C. DeploySnapshot
- D. CreateSnapshot

Answer: D

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/ApiReference-cmd-CreateSnapshot.html>

QUESTION 42

What is an isolated database environment running in the cloud (Amazon RDS) called?

- A. DB Instance

- B. DB Server
- C. DB Unit
- D. DB Volume

Answer: A

QUESTION 43

While signing in REST/ Query requests, for additional security, you should transmit your requests using Secure Sockets Layer (SSL) by using _____.

- A. HTTP
- B. Internet Protocol Security(IPsec)
- C. TLS (Transport Layer Security)
- D. HTTPS

Answer: D

QUESTION 44

What happens to the I/O operations while you take a database snapshot?

- A. I/O operations to the database are suspended for a few minutes while the backup is in progress.
- B. I/O operations to the database are sent to a Replica (if available) for a few minutes while the backup is in progress.
- C. I/O operations will be functioning normally.
- D. I/O operations to the database are suspended for an hour while the backup is in progress.

Answer: A

Explanation:

Creating this DB snapshot on a Single-AZ DB instance results in a brief I/O suspension that typically lasting no more than a few minutes. Multi-AZ DB instances are not affected by this I/O suspension since the backup is taken on the standby.

QUESTION 45

Read Replicas require a transactional storage engine and are only supported for the _____ storage engine.

- A. OracleISAM
- B. MSSQLDB
- C. InnoDB
- D. MyISAM

Answer: C

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 46

When running my DB Instance as a Multi-AZ deployment, can I use the standby for read or write operations?

- A. Yes
- B. Only with MSSQL based RDS
- C. Only for Oracle RDS instances
- D. No

Answer: D

QUESTION 47

When should I choose Provisioned IOPS over Standard RDS storage?

- A. If you have batch-oriented workloads
- B. If you use production online transaction processing (OLTP) workloads
- C. If you have workloads that are not sensitive to consistent performance

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

Amazon RDS provisions that IOPS rate and storage for the lifetime of the DB instance or until you change it. Provisioned IOPS storage is optimized for I/O intensive, online transaction processing (OLTP) workloads that have consistent performance requirements. Provisioned IOPS helps performance tuning.

QUESTION 48

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send ____ minute metrics to Amazon CloudWatch.

- A. 3
- B. 1
- C. 5
- D. 2

Answer: B

QUESTION 49

What is the minimum charge for the data transferred between Amazon RDS and Amazon EC2 Instances in the same Availability Zone?

- A. USD 0.10 per GB
- B. No charge. It is free.
- C. USD 0.02 per GB
- D. USD 0.01 per GB

Answer: B

QUESTION 50

Are Reserved Instances available for Multi-AZ Deployments?

- A. Only for Cluster Compute instances
- B. Yes for all instance types
- C. Only for M3 instance types
- D. No

Answer: B

Explanation:

<https://aws.amazon.com/rds/faqs/>

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 51 – Question 100**

QUESTION 51

Which service enables AWS customers to manage users and permissions in AWS?

- A. AWS Access Control Service (ACS)
- B. AWS Identity and Access Management (IAM)
- C. AWS Identity Manager (AIM)

Answer: B

QUESTION 52

Which Amazon Storage behaves like raw, unformatted, external block devices that you can attach to your instances?

- A. None of these.
- B. Amazon Instance Storage
- C. Amazon EBS
- D. All of these

Answer: C

QUESTION 53

Which Amazon service can I use to define a virtual network that closely resembles a traditional data center?

- A. Amazon VPC
- B. Amazon ServiceBus
- C. Amazon EMR
- D. Amazon RDS

Answer: A

QUESTION 54

What is the command line instruction for running the remote desktop client in Windows?

- A. desk.cpl
- B. mstsc

Answer: B

QUESTION 55

Amazon RDS automated backups and DB Snapshots are currently supported for only the ____ storage engine.

- A. MyISAM
- B. InnoDB

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_MySQL.html

QUESTION 56

MySQL installations default to port ____.

- A. 3306
- B. 443
- C. 80
- D. 1158

Answer: A

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ConnectToInstance.html

QUESTION 57

If you have chosen Multi-AZ deployment, in the event of a planned or unplanned outage of your primary DB Instance, Amazon RDS automatically switches to the standby replica. The automatic failover mechanism simply changes the ____ record of the main DB Instance to point to the standby DB Instance.

- A. DNAME
- B. CNAME
- C. TXT
- D. MX

Answer: B

Explanation:

“When failing over, Amazon RDS simply flips the canonical name record (CNAME) for your DB Instance to point at the standby, which is in turn promoted to become the new primary”

<https://aws.amazon.com/rds/faqs/>

QUESTION 58

If I modify a DB Instance or the DB parameter group associated with the instance, should I reboot the instance for the changes to take effect?

- A. No
- B. Yes

Answer: B

QUESTION 59

If I want to run a database in an Amazon instance, which is the most recommended Amazon storage option?

- A. Amazon Instance Storage
- B. Amazon EBS
- C. You can't run a database inside an Amazon instance
- D. Amazon S3

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Storage.html>

QUESTION 60

In regards to IAM you can edit user properties later, but you cannot use the console to change the ____.

- A. user name
- B. password
- C. default group

Answer: A

QUESTION 61

Can I test my DB Instance against a new version before upgrading?

- A. No
- B. Yes
- C. Only in VPC

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_UpgradeDBInstance.Upgrading.html

QUESTION 62

True or False: If you add a tag that has the same key as an existing tag on a DB Instance, the new value overwrites the old value.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

QUESTION 63

Can I use Provisioned IOPS with VPC?

- A. Only Oracle based RDS
- B. No
- C. Only with MSSQL based RDS
- D. Yes for all RDS instances

Answer: D

QUESTION 64

Making your snapshot public shares all snapshot data with everyone. Can the snapshots with AWS Marketplace product codes be made public?

- A. No
- B. Yes

Answer: B

Explanation:

“Making your snapshot public shares all snapshot data with everyone; however, snapshots with AWS Marketplace product codes cannot be made public. Encrypted snapshots cannot be shared between accounts or made public.”

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html>

QUESTION 65

Fill in the blanks: "To ensure failover capabilities, consider using a ____ for incoming traffic on a network interface".

- A. primary public IP
- B. secondary private IP
- C. secondary public IP
- D. add on secondary IP

Answer: B

Explanation:

To ensure failover capabilities, consider using a secondary private IP for incoming traffic on an elastic network interface. In the event of an instance failure, you can move the interface and/or secondary private IP address to a standby instance

QUESTION 66

If I have multiple Read Replicas for my master DB Instance and I promote one of them, what happens to the rest of the Read Replicas?

- A. The remaining Read Replicas will still replicate from the older master DB Instance
- B. The remaining Read Replicas will be deleted
- C. The remaining Read Replicas will be combined to one read replica

Answer: A

Explanation:

If a source DB instance has several Read Replicas, promoting one of the Read Replicas to a DB instance has no effect on the other replicas.

QUESTION 67

What does Amazon CloudFormation provide?

- A. The ability to setup Autoscaling for Amazon EC2 instances.
- B. None of these.
- C. A templated resource creation for Amazon Web Services.
- D. A template to map network resources for Amazon Web Services.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>

AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you.

QUESTION 68

Can I encrypt connections between my application and my DB Instance using SSL?

- A. No
- B. Yes
- C. Only in VPC
- D. Only in certain regions

Answer: B

QUESTION 69

What are the four levels of AWS Premium Support?

- A. Basic, Developer, Business, Enterprise
- B. Basic, Startup, Business, Enterprise
- C. Free, Bronze, Silver, Gold
- D. All support is free

Answer: A

Explanation:

Q: How are the enhanced AWS Support tiers different from Basic Support?

AWS Basic Support offers all AWS customers access to our Resource Center, Service Health Dashboard, Product FAQs, Discussion Forums, and Support for Health Checks – at no additional charge. Customers who desire a deeper level of support can subscribe to AWS Support at the Developer, Business, or Enterprise level.

<https://aws.amazon.com/premiumsupport/faqs/>

QUESTION 70

What can I access by visiting the URL: <http://status.aws.amazon.com/>?

- A. Amazon Cloud Watch
- B. Status of the Amazon RDS DB
- C. AWS Service Health Dashboard
- D. AWS Cloud Monitor

Answer: C

QUESTION 71

Please select the Amazon EC2 resource which cannot be tagged.

- A. images (AMIs, kernels, RAM disks)
- B. Amazon EBS volumes
- C. Elastic IP addresses
- D. VPCs

Answer: C

Explanation:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html#tag-restrictions

QUESTION 72

Can the string value of 'Key' be prefixed with :aws:"?

- A. Only in GovCloud
- B. Only for S3 not EC2
- C. Yes
- D. No

Answer: D

Explanation:

"The tag key is the required name of the tag. The string value can be from 1 to 128 Unicode characters in length and cannot be prefixed with "aws:" or "rds:"."

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Tagging.html

QUESTION 73

Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production EC2 instances. Which of the following strategies will help prevent a similar situation in the future? The administrator still must be able to:

- launch, start stop, and terminate development resources.
- launch and start production instances.

- A. Create an IAM user, which is not allowed to terminate instances by leveraging production EC2 termination protection.
- B. Leverage resource based tagging along with an IAM user, which can prevent specific users from terminating production EC2 resources.
- C. Leverage EC2 termination protection and multi-factor authentication, which together require users to authenticate before terminating EC2 instances.
- D. Create an IAM user and apply an IAM role which prevents users from terminating production EC2 instances.

Answer: D

QUESTION 74

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

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

Answer: AC

QUESTION 75

Select the incorrect statement

- A. In Amazon EC2, the private IP addresses only returned to Amazon EC2 when the instance is stopped or terminated.
- B. In Amazon VPC, an instance retains its private IP addresses when the instance is stopped.

- C. In Amazon VPC, an instance does NOT retain its private IP addresses when the instance is stopped.
- D. In Amazon EC2, the private IP address is associated exclusively with the instance for its lifetime.

Answer: C

Explanation:

A private IP address remains associated with the network interface when the instance is stopped and restarted, and is released when the instance is terminated.

QUESTION 76

How are the EBS snapshots saved on Amazon S3?

- A. Exponentially
- B. Incrementally
- C. EBS snapshots are not stored in the Amazon S3
- D. Decrementally

Answer: B

QUESTION 77

What is the type of monitoring data (for Amazon EBS volumes) which is available automatically in 5-minute periods at no charge called?

- A. Basic
- B. Primary
- C. Detailed
- D. Local

Answer: A

Explanation:

Monitoring Volumes with CloudWatch

CloudWatch metrics are statistical data that you can use to view, analyze, and set alarms on the operational behavior of your volumes. The following table describes the types of monitoring data available for your Amazon EBS volumes:

Basic

Data is available automatically in 5-minute periods at no charge. This includes data for the root device volumes for EBS-backed instances.

Detailed

Provisioned IOPS SSD (io1) volumes automatically send one-minute metrics to CloudWatch.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/monitoring-volume-status.html>

QUESTION 78

The new DB Instance that is created when you promote a Read Replica retains the backup window period.

- A. TRUE
- B. FALSE

Answer: A

Explanation:

“The new DB instance that is created when you promote a Read Replica retains the backup retention period, backup window period, and parameter group of the former Read Replica source.”

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

QUESTION 79

What happens when you create a topic on Amazon SNS?

- A. The topic is created, and it has the name you specified for it.
- B. An ARN (Amazon Resource Name) is created.
- C. You can create a topic on Amazon SQS, not on Amazon SNS.
- D. This question doesn't make sense.

Answer: B

QUESTION 80

Can I delete a snapshot of the root device of an EBS volume used by a registered AMI?

- A. Only via API
- B. Only via Console
- C. Yes
- D. No

Answer: D

Explanation:

Note that you can't delete a snapshot of the root device of an EBS volume used by a registered AMI. You must first deregister the AMI before you can delete the snapshot.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-snapshot.html>

QUESTION 81

Can I test my DB Instance against a new version before upgrading?

- A. Only in VPC
- B. No
- C. Yes

Answer: C

QUESTION 82

What is the maximum response time for a Business level Premium Support case?

- A. 120 seconds
- B. 1 hour
- C. 10 minutes
- D. 12 hours

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/features/>

QUESTION 83

The ____ service is targeted at organizations with multiple users or systems that use AWS products such as Amazon EC2, Amazon SimpleDB, and the AWS Management Console.

- A. Amazon RDS
- B. AWS Integrity Management
- C. AWS Identity and Access Management
- D. Amazon EMR

Answer: C

Explanation:

https://aws.amazon.com/documentation/iam/?nc1=h_ls

QUESTION 84

True or False: Without IAM, you cannot control the tasks a particular user or system can do and what AWS resources they might use.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

<http://docs.aws.amazon.com/IAM/latest/UserGuide/getting-setup.html>

QUESTION 85

When you use the AWS Management Console to delete an IAM user, IAM also deletes any signing certificates and any access keys belonging to the user.

- A. FALSE
- B. This is configurable
- C. TRUE

Answer: C

Explanation:

When you use the AWS Management Console to delete an IAM user, IAM automatically deletes the following information for you:

The user

Any group memberships -- that is, the user is removed from any IAM groups that the user was a member of:

Any password associated with the user

Any access keys belonging to the user

All inline policies embedded in the user (policies that are applied to a user via group permissions are not affected)

Note!

Any managed policies attached to the user are detached from the user when the user is deleted. Managed policies are not deleted when you delete a user.

Any associated MFA device

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_users_manage.html#id_users_deleting_console

QUESTION 86

When automatic failover occurs, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the _____ to return information about events related to your DB Instance.

- A. FetchFailure
- B. DescribeFailure
- C. DescribeEvents
- D. FetchEvents

Answer: C

Explanation:

Q: Will I be alerted when automatic failover occurs?

Yes, Amazon RDS will emit a DB Instance event to inform you that automatic failover occurred. You can use the DescribeEvents to return information about events related to your DB Instance, or click the "DB Events" section of the AWS Management Console

<https://aws.amazon.com/rds/faqs/>

QUESTION 87

What is the default maximum number of MFA devices in use per AWS account (at the root account level)?

- A. 1
- B. 5
- C. 15
- D. 10

Answer: A

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-limits.html

QUESTION 88

Do the Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. Only if instructed to when created
- B. Yes
- C. No

Answer: B

Explanation:

Data persistence

An EBS volume is off-instance storage that can persist independently from the life of an instance. You continue to pay for the volume usage as long as the data persists.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumes.html#EBSFeatures>

QUESTION 89

Can we attach an EBS volume to more than one EC2 instance at the same time?

- A. Yes.
- B. No
- C. Only EC2-optimized EBS volumes
- D. Only in read mode

Answer: B

Explanation:

EBS is network attached storage that can only be attached to one instance at a time
<https://aws.amazon.com/ebs/getting-started/>

QUESTION 90

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

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

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#default-security-group>

A default security group is named default, and it has an ID assigned by AWS. The following are the initial settings for each default security group:

Allow inbound traffic only from other instances associated with the default security group

Allow all outbound traffic from the instance

The default security group specifies itself as a source security group in its inbound rules. This is what allows instances associated with the default security group to communicate with other instances associated with the default security group.

QUESTION 91

What does Amazon Route53 provide?

- A. A global Content Delivery Network.
- B. None of these.
- C. A scalable Domain Name System.
- D. An SSH endpoint for Amazon EC2.

Answer: C

Explanation:

<https://aws.amazon.com/route53/>

QUESTION 92

What does Amazon ElastiCache provide?

- A. A service by this name doesn't exist. Perhaps you mean Amazon CloudCache.
- B. A virtual server with a huge amount of memory.
- C. A managed In-memory cache service.
- D. An Amazon EC2 instance with the Memcached software already pre-installed.

Answer: C

QUESTION 93

How many Elastic IP by default in Amazon Account?

- A. 1 Elastic IP
- B. 3 Elastic IP

- C. 5 Elastic IP
- D. 0 Elastic IP

Answer: C

Explanation:

“By default, all AWS accounts are limited to 5 Elastic IP addresses, because public (IPv4) Internet addresses are a scarce public resource.”

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

QUESTION 94

What is a Security Group?

- A. None of these.
- B. A list of users that can access Amazon EC2 instances.
- C. An Access Control List (ACL) for AWS resources.
- D. A firewall for inbound traffic, built-in around every Amazon EC2 instance.

Answer: D

Explanation:

A security group acts as a virtual firewall for your instance to control inbound and outbound traffic.

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html

QUESTION 95

The one-time payment for Reserved Instances is _____ refundable if the reservation is cancelled.

- A. always
- B. in some circumstances
- C. never

Answer: C

Explanation:

the one-time fee is non-refundable.

<https://aws.amazon.com/ec2/purchasing-options/reserved-instances/buyer/>

QUESTION 96

Please select the Amazon EC2 resource which can be tagged.

- A. key pairs
- B. Elastic IP addresses
- C. placement groups
- D. Amazon EBS snapshots

Answer: D

Explanation:

Placement group and Elastic IP cannot be tagged.

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

Snapshots can be tagged: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

QUESTION 97

If an Amazon EBS volume is the root device of an instance, can I detach it without stopping the instance?

- A. Yes but only if Windows instance
- B. No
- C. Yes
- D. Yes but only if a Linux instance

Answer: B

Explanation:

“If an EBS volume is the root device of an instance, you must stop the instance before you can detach the volume.”

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

QUESTION 98

If you are using Amazon RDS Provisioned IOPS storage with MySQL and Oracle database engines, you can scale the throughput of your database Instance by specifying the IOPS rate from _____.

- A. 1,000 to 1,00,000
- B. 100 to 1,000
- C. 10,000 to 1,00,000
- D. 1,000 to 10,000

Answer: D

Explanation:

If you are using RDS Provisioned IOPS, you can also scale the throughput of your DB Instance by specifying the IOPS rate from 1,000 IOPS to 10,000 IOPS in 1,000 IOPS

<https://aws.amazon.com/rds/mysql/>

QUESTION 99

Every user you create in the IAM system starts with _____.

- A. full permissions
- B. no permissions
- C. partial permissions

Answer: B

Explanation:

Permissions let you specify who has access to AWS resources, and what actions they can perform on those resources. Every IAM user starts with no permissions.

http://docs.aws.amazon.com/IAM/latest/UserGuide/access_permissions.html#NoDefaultPermissions

QUESTION 100

Amazon RDS DB snapshots and automated backups are stored in _____.

- A. Amazon S3
- B. Amazon ECS Volume
- C. Amazon RDS
- D. Amazon EMR

Answer: A

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 101 – Question 150**

QUESTION 101

What is the maximum key length of a tag?

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

Answer: D

QUESTION 102

Groups can't _____.

- A. be nested more than 3 levels
- B. be nested at all
- C. be nested more than 4 levels
- D. be nested more than 2 levels

Answer: B

QUESTION 103

Amazon EC2 has no Amazon Resource Names (ARNs) because you can't specify a particular Amazon EC2 resource in an IAM policy.

- A. TRUE
- B. FALSE

Answer: B

Explanation:

<http://blogs.aws.amazon.com/security/post/Tx29HCT3ABL7LP3/Resource-level-Permissions-for-EC2-Controlling-Management-Access-on-Specific-Ins>

QUESTION 104

Can I initiate a "forced failover" for my MySQL Multi-AZ DB Instance deployment?

- A. Only in certain regions
- B. Only in VPC
- C. Yes
- D. No

Answer: C

Explanation:

If your DB instance is a Multi-AZ deployment, you can force a failover from one availability zone to another when you select the Reboot option. When you force a failover of your DB instance, Amazon RDS automatically switches to a standby replica in another Availability Zone and updates the DNS record for the DB instance to point to the standby DB instance. As a result, you will need to clean up and re-establish any existing connections to your DB instance. Reboot with failover is beneficial when you want to simulate a failure of a DB instance for testing, or restore operations to the original AZ after a failover occurs.

Source: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RebootInstance.html

QUESTION 105

A group can contain many users. Can a user belong to multiple groups?

- A. Yes always
- B. No
- C. Yes but only if they are using two factor authentication
- D. Yes but only in VPC

Answer: A

Explanation:

A group can contain many users, and a user can belong to multiple groups.

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups.html

QUESTION 106

Is the encryption of connections between my application and my DB Instance using SSL for the MySQL server engines available?

- A. Yes
- B. Only in VPC
- C. Only in certain regions
- D. No

Answer: A

Explanation:

<https://aws.amazon.com/rds/faqs/>

QUESTION 107

Which AWS instance address has the following characteristics?

"If you stop an instance, its Elastic IP address is unmapped, and you must remap it when you restart the instance."

- A. VPC Addresses
- B. EC2 Addresses
- C. Both A and B
- D. None of these

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/elastic-ip-addresses-eip.html>

Stopping an instance

EC2-Classical

If you stop an instance, its Elastic IP address is disassociated, and you must re-associate the Elastic IP address when you restart the instance.

EC2-VPC

If you stop an instance, its Elastic IP address remains associated.

QUESTION 108

True or False: Common points of failures like generators and cooling equipment are shared across Availability Zones.

- A. TRUE
- B. FALSE

Answer: B

QUESTION 109

Please select the most correct answer regarding the persistence of the Amazon Instance Store.

- A. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance
- B. The data on an instance store volume is lost when the security group rule of the associated instance is changed
- C. The data on an instance store volume persists even after associated Amazon EC2 instance is deleted

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Storage.html>

QUESTION 110

Multi-AZ deployment _____ supported for Microsoft SQL Server DB Instances.

- A. is not currently
- B. is as of 2013
- C. is planned to be in 2014
- D. will never be

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2014/05/19/amazon-rds-for-sqlserver-introduces-multi-az-support/>

QUESTION 111

Security groups act like a firewall at the instance level, whereas _____ are an additional layer of security that act at the subnet level.

- A. DB Security Groups
- B. VPC Security Groups
- C. network ACLs

Answer: C

QUESTION 112

What does Amazon Elastic Beanstalk provide?

- A. An application container on top of Amazon Web Services.
- B. A scalable storage appliance on top of Amazon Web Services.
- C. A scalable cluster of EC2 instances.
- D. A service by this name doesn't exist.

Answer: A

QUESTION 113

Is the SQL Server Audit feature supported in the Amazon RDS SQL Server engine?

- A. No
- B. Yes

Answer: A

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_SQLServer.html

QUESTION 114

Are you able to integrate a multi-factor token service with the AWS Platform?

- A. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platform.
- B. No, you cannot integrate multi-factor token devices with the AWS platform.
- C. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platform.

Answer: C

Explanation:

Private MFA does not apply here.

Q. What is AWS MFA?

AWS multi-factor authentication (AWS MFA) provides an extra level of security that you can apply to your AWS environment. You can enable AWS MFA for your AWS account and for individual AWS Identity and Access Management (IAM) users you create under your account.

QUESTION 115

My Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do?

- A. You will need to delete the Read Replica and create a new one to replace it.
- B. You will need to disassociate the DB Engine and re associate it.
- C. The instance should be deployed to Single AZ and then moved to Multi-AZ once again.
- D. You will need to delete the DB Instance and create a new one to replace it.

Answer: A

Explanation:

Q: My Amazon RDS for MySQL Read Replica appears "stuck" after a Multi-AZ failover and is unable to obtain or apply updates from the source DB Instance. What do I do? ... To resolve the current issue, you will need to delete the Read Replica and create a new one to replace it. "

<https://aws.amazon.com/rds/faqs/>

QUESTION 116

Which DNS name can only be resolved within Amazon EC2?

- A. Internal DNS name
- B. External DNS name
- C. Global DNS name
- D. Private DNS name

Answer: D

Explanation:

Using Private Hosted Zones

If you want to access the resources in your VPC using custom DNS domain names, such as example.com, instead of using private IP addresses or AWS-provided private DNS hostnames, you can create a private hosted zone in Amazon Route 53. A private hosted zone is a container that holds information about how you want to route traffic for a domain and its subdomains within one or more VPCs without exposing your resources to the Internet. You can then create Amazon Route 53 resource record sets, which determine how Amazon Route 53 responds to queries for your domain and subdomains. For example, if you want browser requests for example.com to be routed to a web server in your VPC, you'll create an A record in your private hosted zone and specify the IP address of that web server. For more information about creating a private hosted zone, see Working with Private Hosted Zones in the Amazon Route 53 Developer Guide.

QUESTION 117

If your DB instance runs out of storage space or file system resources, its status will change to ____ and your DB Instance will no longer be available.

- A. storage-overflow
- B. storage-full
- C. storage-exceed
- D. storage-overage

Answer: B

Explanation:

<https://aws.amazon.com/ko/premiumsupport/knowledge-center/rds-out-of-storage/>

QUESTION 118

Is it possible to access your EBS snapshots?

- A. Yes, through the Amazon S3 APIs.
- B. Yes, through the Amazon EC2 APIs.
- C. No, EBS snapshots cannot be accessed; they can only be used to create a new EBS volume.
- D. EBS doesn't provide snapshots.

Answer: B

Explanation:

https://aws.amazon.com/ebs/faqs/?nc1=h_ls

Q: Will I be able to access my snapshots using the regular Amazon S3 API?

No, snapshots are only available through the Amazon EC2 API.

QUESTION 119

Does Amazon RDS for SQL Server currently support importing data into the msdb database?

- A. No
- B. Yes

Answer: A

Explanation:

Amazon RDS for SQL Server does not currently support importing data into the msdb database, though we do support SQL Server Agent jobs. Some SQL Server features that use the msdb database, such as Database Mail and Replication, are not currently supported in Amazon RDS.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html#SQLServer.Procedural.Importing.Procedure>

QUESTION 120

Does Route 53 support MX Records?

- A. Yes.
- B. It supports CNAME records, but not MX records.
- C. No.
- D. Only Primary MX records. Secondary MX records are not supported.

Answer: A

Explanation:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/ResourceRecordTypes.html#MXFormat>

QUESTION 121

Because of the extensibility limitations of striped storage attached to Windows Server, Amazon RDS does not currently support increasing storage on a ____ DB Instance.

- A. SQL Server
- B. MySQL
- C. Oracle

Answer: A

QUESTION 122

Which Amazon storage do you think is the best for my database-style applications that frequently encounter many random reads and writes across the dataset?

- A. None of these.
- B. Amazon Instance Storage
- C. Any of these
- D. Amazon EBS

Answer: D

Explanation:

“Amazon EBS is particularly helpful for database-style applications that frequently encounter many random reads and writes across the data set.”

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html>

QUESTION 123

Select the correct set of steps for exposing the snapshot only to specific AWS accounts.

- A. Select public for all the accounts and check mark those accounts with whom you want to expose the snapshots and click Save.

- B. Select Private, enter the IDs of those AWS accounts, and click Save.
- C. Select Public, enter the IDs of those AWS accounts, and click Save.
- D. Select Public, mark the IDs of those AWS accounts as private, and click Save.

Answer: B

Explanation:

“To expose the snapshot to only specific AWS accounts, choose Private, enter the ID of the AWS account (without hyphens) in the AWS Account Number field, and choose Add Permission. Repeat until you’ve added all the required AWS accounts.”

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modifying-snapshot-permissions.html>

QUESTION 124

Is decreasing the storage size of a DB Instance permitted?

- A. Depends on the RDMS used
- B. Yes
- C. No

Answer: C

Explanation:

“note that you cannot reduce storage size once it has been allocated”

Source: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuideCHAP_Storage.html#CHAP_Storage.FactsAbout

QUESTION 125

When should I choose Provisioned IOPS over Standard RDS storage?

- A. If you use production online transaction processing (OLTP) workloads
- B. If you have batch-oriented workloads
- C. If you have workloads that are not sensitive to consistent performance

Answer: B

QUESTION 126

In the context of MySQL, version numbers are organized as MySQL version = X.Y.Z. What does X denote here?

- A. release level
- B. minor version
- C. version number
- D. major version

Answer: D

QUESTION 127

In the 'Detailed' monitoring data available for your Amazon EBS volumes, Provisioned IOPS volumes automatically send ____ minute metrics to Amazon CloudWatch.

- A. 5
- B. 2
- C. 1
- D. 3

Answer: C

QUESTION 128

It is advised that you watch the Amazon CloudWatch "____" metric (available via the AWS Management Console or Amazon Cloud Watch APIs) carefully and recreate the Read Replica should it fall behind due to replication errors.

- A. Write Lag
- B. Read Replica
- C. Replica Lag
- D. Single Replica

Answer: C

Explanation:

The amount of time a Read Replica DB instance lags behind the source DB instance. Applies to MySQL, MariaDB, and PostgreSQL Read Replicas.

<http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/rds-metricscollected.html>

QUESTION 129

Can the string value of 'Key' be prefixed with laws?

- A. No
- B. Only for EC2 not S3
- C. Yes
- D. Only for S3 not EC

Answer: A

Explanation:

<http://docs.aws.amazon.com/cli/latest/reference/rds/list-tags-for-resource.html>

QUESTION 130

By default what are ENIs that are automatically created and attached to instances using the EC2 console set to do when the attached instance terminates?

- A. Remain as is
- B. Terminate
- C. Hibernate
- D. Pause

Answer: B

Explanation:

By default, elastic network interfaces that are automatically created and attached to instances using the console are set to terminate when the instance terminates. However, network interfaces created using the command line interface aren't set to terminate when the instance terminates.

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

QUESTION 131

Are you able to integrate a multi-factor token service with the AWS Platform?

- A. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platform.
- B. No, you cannot integrate multi-factor token devices with the AWS platform.
- C. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platform.

Answer: C

QUESTION 132

You can use ____ and ____ to help secure the instances in your VPC.

- A. security groups and multi-factor authentication
- B. security groups and 2-Factor authentication
- C. security groups and biometric authentication
- D. security groups and network ACLs

Answer: D

QUESTION 133

Fill in the blanks: ____ is a durable, block-level storage volume that you can attach to a single, running Amazon EC2 instance.

- A. Amazon S3
- B. Amazon EBS
- C. None of these
- D. All of these

Answer: B

QUESTION 134

Do the Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance?

- A. No
- B. Only if instructed to when created
- C. Yes

Answer: C

QUESTION 135

If I want my instance to run on a single-tenant hardware, which value do I have to set the instance's tenancy attribute to?

- A. dedicated
- B. isolated
- C. one
- D. reserved

Answer: A

Explanation:

<http://aws.amazon.com/ec2/dedicated-hosts/>

QUESTION 136

What does Amazon RDS stand for?

- A. Regional Data Server.
- B. Relational Database Service.
- C. Nothing.
- D. Regional Database Service.

Answer: B

QUESTION 137

What is the maximum response time for a Business level Premium Support case?

- A. 30 minutes
- B. You always get instant responses (within a few seconds)
- C. 10 minutes
- D. 1 hour

Answer: D

QUESTION 138

What does Amazon ELB stand for?

- A. Elastic Linux Box.
- B. Encrypted Linux Box.
- C. Encrypted Load Balancing.
- D. Elastic Load Balancing.

Answer: D

QUESTION 139

What does Amazon CloudFormation provide?

- A. None of these.
- B. The ability to setup Autoscaling for Amazon EC2 instances.
- C. A template to map network resources for Amazon Web Services.
- D. A templated resource creation for Amazon Web Services.

Answer: C

QUESTION 140

Is there a limit to the number of groups you can have?

- A. Yes for all users except root
- B. No
- C. Yes unless special permission granted
- D. Yes for all users

Answer: C

Explanation:

Currently you can request to increase the limit on users per AWS account, groups per AWS account, roles per AWS account, instance profiles per AWS account, and server certificates per AWS account.

http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_iam-limits.html

QUESTION 141

Location of Instances are ____.

- A. Regional
- B. based on Availability Zone
- C. Global

Answer: B

Explanation:

Regions and Availability Zones

Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each region is a separate geographic area. Each region has multiple, isolated locations known as Availability Zones. Amazon EC2 provides you the ability to place resources, such as instances, and data in multiple locations. Resources aren't replicated across regions unless you do so specifically.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions-availability-zones>

QUESTION 142

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

- A. You can create an encrypted tunnel to VPC, but you don't own the connection.
- B. Yes, it's called Amazon Dedicated Connection.
- C. No, AWS only allows access from the public Internet.
- D. Yes, it's called Direct Connect.

Answer: D

QUESTION 143

What is the maximum response time for a Business level Premium Support case?

- A. 30 minutes
- B. 1 hour
- C. 12 hours
- D. 10 minutes

Answer: B

QUESTION 144

Does Dynamic DB support in-place atomic updates?

- A. It is not defined
- B. No
- C. Yes
- D. It does support in-place non-atomic updates

Answer: C

Explanation:

Q: Does DynamoDB support in-place atomic updates?

Amazon DynamoDB supports fast in-place updates. You can increment or decrement a numeric attribute in a row using a single API call. Similarly, you can atomically add or remove to sets, lists, or maps.

<https://aws.amazon.com/dynamodb/faqs/>

QUESTION 145

Is there a method in the IAM system to allow or deny access to a specific instance?

- A. Only for VPC based instances
- B. Yes
- C. No

Answer: B

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html#policy-eval-denyallow

- 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 overrides this default.

- An explicit deny overrides any allows.

QUESTION 146

What is an isolated database environment running in the cloud (Amazon RDS) called?

- A. DB Instance
- B. DB Unit
- C. DB Server
- D. DB Volume

Answer: A

QUESTION 147

What does Amazon SES stand for?

- A. Simple Elastic Server
- B. Simple Email Service
- C. Software Email Solution
- D. Software Enabled Server

Answer: B

Explanation:

<http://aws.amazon.com/ses/>

QUESTION 148

Amazon S3 doesn't automatically give a user who creates _____ permission to perform other actions on that bucket or object.

- A. a file
- B. a bucket or object
- C. a bucket or file
- D. a object or file

Answer: B

Explanation:

Amazon S3 doesn't automatically give a user who creates a bucket or object permission to perform other actions on that bucket or object. Therefore, in your IAM policies, you must explicitly give users permission to use the Amazon S3 resources they create.

http://docs.aws.amazon.com/IAM/latest/UserGuide/IAM_UseCases.html

QUESTION 149

While creating the snapshots using the command line tools, which command should I be using?

- A. ec2-deploy-snapshot
- B. ec2-fresh-snapshot
- C. ec2-create-snapshot
- D. ec2-new-snapshot

Answer: C

QUESTION 150

Typically, you want your application to check whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an ____ node in the response from the Amazon RDS API.

- A. Incorrect
- B. Error
- C. FALSE

Answer: B

Explanation:

Typically, you want your application to check whether a request generated an error before you spend any time processing results. The easiest way to find out if an error occurred is to look for an Error node in the response from the Amazon RDS API.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/APITroubleshooting.html>

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 151 – Question 200**

QUESTION 151

Can I detach the primary (eth0) network interface when the instance is running or stopped?

- A. Yes, You can
- B. No. You cannot
- C. Depends on the state of the interface at the time

Answer: B

Explanation:

Each instance in a VPC has a default elastic network interface (the primary network interface, eth0) that is assigned a private IP address from the IP address range of your VPC. You cannot detach a primary network interface from an instance.

QUESTION 152

What's an ECU?

- A. Extended Cluster User.
- B. None of these.
- C. Elastic Computer Usage.
- D. Elastic Compute Unit.

Answer: B

Explanation:

The EC2 Compute Unit (ECU) provides the relative measure of the integer processing power of an Amazon EC2 instance. <https://aws.amazon.com/ec2/faqs/>

QUESTION 153

REST or Query requests are HTTP or HTTPS requests that use an HTTP verb (such as GET or POST) and a parameter named Action or Operation that specifies the API you are calling.

- A. FALSE
- B. TRUE

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/APIReference/Query-Requests.html>

QUESTION 154

What is the charge for the data transfer incurred in replicating data between your primary and standby?

- A. No charge. It is free.
- B. Double the standard data transfer charge.
- C. Same as the standard data transfer charge.

D. Half of the standard data transfer charge.

Answer: A

Explanation:

Q: How much do Read Replicas cost? When does billing begin and end?

A Read Replica is billed as a standard DB Instance and at the same rates. Click here for more information on DB Instance billing visit this FAQ. Just like a standard DB Instance, the rate per “DB Instance hour” for a Read Replica is determined by the DB Instance class of the Read Replica – please see Amazon RDS detail page for up-to-date pricing. You are not charged for the data transfer incurred in replicating data between your source DB Instance and Read Replica. Billing for a Read Replica begins as soon as the Read Replica has been successfully created (i.e. when status is listed as “active”). The Read Replica will continue being billed at standard Amazon RDS DB Instance hour rates until you issue a command to delete it.

QUESTION 155

Does AWS Direct Connect allow you access to all Availabilities Zones within a Region?

- A. Depends on the type of connection
- B. No
- C. Yes
- D. Only when there's just one availability zone in a region. If there are more than one, only one availability zone can be accessed directly

Answer: C

QUESTION 156

What does the "Server Side Encryption" option on Amazon S3 provide?

- A. It provides an encrypted virtual disk in the Cloud.
- B. It doesn't exist for Amazon S3, but only for Amazon EC2.
- C. It encrypts the files that you send to Amazon S3, on the server side.
- D. It allows to upload files using an SSL endpoint, for a secure transfer.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingServerSideEncryption.html>

Server-side encryption is about protecting data at rest. Server-side encryption with Amazon S3-managed encryption keys (SSE-S3) employs strong multi-factor encryption.

Amazon S3 encrypts each object with a unique key. As an additional safeguard, it encrypts the key itself with a master key that it regularly rotates. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data.

QUESTION 157

What does Amazon EBS stand for?

- A. Elastic Block Storage
- B. Elastic Business Server
- C. Elastic Blade Server
- D. Elastic Block Store

Answer: D

Explanation:

<https://aws.amazon.com/ebs/>

Amazon Elastic Block Store (EBS)

Amazon Elastic Block Store (Amazon EBS) provides persistent block level storage volumes for use with Amazon EC2 instances in the AWS Cloud. Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure, offering high availability and durability. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes – all while paying a low price for only what you provision.

QUESTION 158

Within the IAM service a GROUP is regarded as a:

- A. A collection of AWS accounts.
- B. It's the group of EC2 machines that gain the permissions specified in the GROUP.
- C. There's no GROUP in IAM, but only USERS and RESOURCES.
- D. A collection of users.

Answer: D

Explanation:

Use groups to assign permissions to IAM users

Instead of defining permissions for individual IAM users, it's usually more convenient to create groups that relate to job functions (administrators, developers, accounting, etc.), define the relevant permissions for each group, and then assign IAM users to those groups. All the users in an IAM group inherit the permissions assigned to the group. That way, you can make changes for everyone in a group in just one place. As people move around in your company, you can simply change what IAM group their IAM user belongs to.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-groups-for-permissions>

QUESTION 159

A ____ is the concept of allowing (or disallowing) an entity such as a user, group, or role some type of access to one or more resources.

- A. user
- B. AWS Account
- C. resource
- D. permission

Answer: D

Explanation:

A permission is the concept of allowing (or disallowing) an entity such as a user, group, or role some type of access to one or more resources.

QUESTION 160

After an Amazon VPC instance is launched, can I change the VPC security groups it belongs to?

- A. No. You cannot.
- B. Yes. You can.
- C. Only if you are the root user.
- D. Only if the tag "VPC_Change_Group" is true.

Answer: B

Explanation:

Security groups are associated with network interfaces. After you launch an instance, you can change the security groups associated with the instance, which changes the security groups associated with the primary network interface (eth0).

QUESTION 161

Do the system resources on the Micro instance meet the recommended configuration for Oracle?

- A. Yes completely
- B. Yes but only for certain situations
- C. Not in any circumstance

Answer: B

Explanation:

We recommend that you use db.t1.micro instances with Oracle to test setup and connectivity only; the system resources for a db.t1.micro instance do not meet the recommended configuration for Oracle. No Oracle options are supported on a db.t1.micro instance.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuideConcepts.DBInstanceClass.html#Concepts.DBInstanceClass.Previous>

QUESTION 162

Will I be charged if the DB instance is idle?

- A. No
- B. Yes

- C. Only is running in GovCloud
- D. Only if running in VPC

Answer: B

QUESTION 163

To help you manage your Amazon EC2 instances, images, and other Amazon EC2 resources, you can assign your own metadata to each resource in the form of _____.

- A. special filters
- B. functions
- C. tags
- D. wildcards

Answer: C

QUESTION 164

Are you able to integrate a multi-factor token service with the AWS Platform?

- A. No, you cannot integrate multi-factor token devices with the AWS platform.
- B. Yes, you can integrate private multi-factor token devices to authenticate users to the AWS platform.
- C. Yes, using the AWS multi-factor token devices to authenticate users on the AWS platform.

Answer: C

QUESTION 165

True or False: When you add a rule to a DB security group, you do not need to specify port number or protocol.

- A. Depends on the RDMS used
- B. TRUE
- C. FALSE

Answer: B

Explanation:

You no need to specify Port or Protocol for DB security group. Only VPC/instance security group only needed.
<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.RDSSecurityGroups.html>

QUESTION 166

Is there a limit to the number of groups you can have?

- A. Yes for all users
- B. Yes for all users except root
- C. No
- D. Yes unless special permission granted

Answer: A

Explanation:

<http://aws.amazon.com/rds/faqs/#46>

QUESTION 167

Can I initiate a "forced failover" for my Oracle Multi-AZ DB Instance deployment?

- A. Yes
- B. Only in certain regions
- C. Only in VPC
- D. No

Answer: A

Explanation:

<https://aws.amazon.com/public-data-sets/>

QUESTION 168

Amazon EC2 provides a repository of public data sets that can be seamlessly integrated into AWS cloud-based applications. What is the monthly charge for using the public data sets?

- A. A 1 time charge of 10\$ for all the datasets
- B. 1\$ per dataset per month
- C. 10\$ per month for all the datasets
- D. There is no charge for using the public data sets

Answer: D

QUESTION 169

In the Amazon RDS Oracle DB engine, the Database Diagnostic Pack and the Database Tuning Pack are only available with ____.

- A. Oracle Standard Edition
- B. Oracle Express Edition
- C. Oracle Enterprise Edition
- D. None of these

Answer: C

Explanation:

<https://www.pythian.com/blog/a-most-simple-cloud-is-amazon-rds-for-oracle-right-for-you/>

QUESTION 170

Without ____, you must either create multiple AWS accounts-each with its own billing and subscriptions to AWS products- or your employees must share the security credentials of a single AWS account.

- A. Amazon RDS
- B. Amazon Glacier
- C. Amazon EMR
- D. Amazon IAM

Answer: D

QUESTION 171

Amazon RDS supports SOAP only through ____.

- A. HTTP or HTTPS
- B. TCP/IP
- C. HTTP
- D. HTTPS

Answer: D

Explanation:

Amazon RDS supports SOAP only through HTTPS

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/using-soap-api.html>

QUESTION 172

The Amazon EC2 web service can be accessed using the ____ web services messaging protocol. This interface is described by a Web Services Description Language (WSDL) document.

- A. SOAP
- B. DCOM
- C. CORBA
- D. XML-RPC

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSECommerceService/latest/DG/WSDLLocation.html>

QUESTION 173

Is creating a Read Replica of another Read Replica supported?

- A. Only in VPC
- B. Yes
- C. Only in certain regions
- D. No

Answer: B

QUESTION 174

What is the charge for the data transfer incurred in replicating data between your primary and standby?

- A. Same as the standard data transfer charge
- B. Double the standard data transfer charge
- C. No charge. It is free
- D. Half of the standard data transfer charge

Answer: C

QUESTION 175

HTTP Query-based requests are HTTP requests that use the HTTP verb GET or POST and a Query parameter named ____.

- A. Action
- B. Value
- C. Reset
- D. Retrieve

Answer: A

Explanation:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/using-with-s3-actions.html>

QUESTION 176

What happens to the I/O operations while you take a database snapshot?

- A. I/O operations to the database are suspended for an hour while the backup is in progress.
- B. I/O operations to the database are sent to a Replica (if available) for a few minutes while the backup is in progress.
- C. I/O operations will be functioning normally.
- D. I/O operations to the database are suspended for a few minutes while the backup is in progress.

Answer: D

QUESTION 177

Amazon RDS creates an SSL certificate and installs the certificate on the DB Instance when Amazon RDS provisions the instance. These certificates are signed by a certificate authority. The ____ is stored at <https://rds.amazonaws.com/doc/rds-ssl-ca-cert.pem>.

- A. private key
- B. foreign key
- C. public key
- D. protected key

Answer: C

Explanation:

<https://aws.amazon.com/blogs/aws/amazon-rds-sql-server-ssl-support/>

QUESTION 178

____ embodies the "share-nothing" architecture and essentially involves breaking a large database into several smaller databases. Common ways to split a database include 1) splitting tables that are not joined in the same query onto different hosts or 2) duplicating a table across multiple hosts and then using a hashing algorithm to determine which host receives a given update.

- A. Sharding
- B. Failure recovery
- C. Federation
- D. DDL operations

Answer: A

Explanation:

Sharding embodies the “share-nothing” architecture and essentially just involves breaking a larger database up into smaller databases. Common ways to split a database are:

Splitting tables that are not joined in the same query onto different hosts.

Duplicating a table across multiple hosts and then splitting where a row goes.

More detailed information on the pros and cons of sharding can be found at the following sites:

<http://technoroy.blogspot.com/2008/07/shard-database-design.html>

<http://www.hibernate.org/subprojects/shards.html>

How Amazon RDS Helps With Sharding Maintenance Overhead

QUESTION 179

What is the name of licensing model in which I can use your existing Oracle Database licenses to run Oracle deployments on Amazon RDS?

- A. Bring Your Own License
- B. Role Bases License
- C. Enterprise License
- D. License Included

Answer: A

Explanation:

<https://aws.amazon.com/oracle/>

QUESTION 180

When you resize the Amazon RDS DB instance, Amazon RDS will perform the upgrade during the next maintenance window. If you want the upgrade to be performed now, rather than waiting for the maintenance window, specify the ____ option.

- A. ApplyNow
- B. ApplySoon
- C. ApplyThis
- D. ApplyImmediately

Answer: D

Explanation:

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.DBInstance.Modifying.html>

QUESTION 181

Does Amazon Route 53 support NS Records?

- A. Yes, it supports Name Service records.
- B. No.
- C. It supports only MX records.
- D. Yes, it supports Name Server records.

Answer: D

Explanation:

<https://aws.amazon.com/route53/faqs/>

QUESTION 182

The SQL Server ____ feature is an efficient means of copying data from a source database to your DB Instance. It writes the data that you specify to a data file, such as an ASCII file.

- A. bulk copy
- B. group copy
- C. dual copy

D. mass copy

Answer: A

Explanation:

The SQL Server bulk copy feature is an efficient means of copying data from a source database to your DB Instance. Bulk copy writes the data that you specify to a data file, such as an ASCII file. You can then run bulk copy again to write the contents of the file to the destination DB Instance.

<http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/SQLServer.Procedural.Importing.html>

QUESTION 183

When using consolidated billing there are two account types. What are they?

- A. Paying account and Linked account.
- B. Parent account and Child account.
- C. Main account and Sub account.
- D. Main account and Secondary account.

Answer: A

Explanation:

You sign up for Consolidated Billing in the AWS Billing and Cost Management console, and designate your account as a payer account. Now your account can pay the charges of the other accounts, which are called linked accounts. The payer account and the accounts linked to it are called a Consolidated Billing account family.

Source: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html>

QUESTION 184

A ____ is a document that provides a formal statement of one or more permissions.

- A. policy
- B. permission
- C. Role
- D. resource

Answer: A

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

QUESTION 185

In the Amazon RDS which uses the SQL Server engine, what is the maximum size for a Microsoft SQL Server DB Instance with SQL Server Express edition?

- A. 10 GB per DB
- B. 100 GB per DB
- C. 2 TB per DB
- D. 1TB per DB

Answer: A

Explanation:

The maximum storage size for a Microsoft SQL Server DB Instance is 4 TB for all instances except the SQL Server Express edition, which limits storage to a total of 300 GB. The minimum storage size for a Microsoft SQL Server DB Instance is 20 GB for the Microsoft SQL Server Express and Web Editions and 200 GB for the Standard and Enterprise Editions.

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_SQLServer.html

QUESTION 186

Regarding the attaching of ENI to an instance, what does 'warm attach' refer to?

- A. Attaching an ENI to an instance when it is stopped.
- B. This question doesn't make sense.
- C. Attaching an ENI to an instance when it is running.
- D. Attaching an ENI to an instance during the launch process.

Answer: A

Explanation:

Best Practices for Configuring Elastic Network Interfaces

You can attach an elastic network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach).

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#best-practices-for-configuring-network-interfaces>

QUESTION 187

If I scale the storage capacity provisioned to my DB Instance by mid of a billing month, how will I be charged?

- A. You will be charged for the highest storage capacity you have used
- B. On a proration basis
- C. You will be charged for the lowest storage capacity you have used

Answer: B

Explanation:

<https://aws.amazon.com/ebs/pricing/>

QUESTION 188

You can modify the backup retention period; valid values are 0 (for no backup retention) to a maximum of ____ days.

- A. 45
- B. 35
- C. 15
- D. 5

Answer: B

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithAutomatedBackups.html

QUESTION 189

A Provisioned IOPS volume must be at least ____ GB in size.

- A. 1
- B. 50
- C. 20
- D. 10

Answer: D

Explanation:

<https://aws.amazon.com/ebs/details/>

QUESTION 190

Will I be alerted when automatic failover occurs?

- A. Only if SNS configured
- B. No
- C. Yes
- D. Only if Cloudwatch configured

Answer: A

Explanation:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_Events.html#USER_Events.Messages

We only get alert on failover if we configure the SNS Notification.

QUESTION 191

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

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

Answer: C

Explanation:

While EBS volumes are locked to the Availability Zone in which they reside, snapshots are available throughout their region of residence).

QUESTION 192

If you're unable to connect via SSH to your EC2 instance, which of the following should you check and possibly correct to restore connectivity?

- A. Adjust Security Group to permit egress traffic over TCP port 443 from your IP.
- B. Configure the IAM role to permit changes to security group settings.
- C. Modify the instance security group to allow ingress of ICMP packets from your IP.
- D. Adjust the instance's Security Group to permit ingress traffic over port 22 from your IP.
- E. Apply the most recently released Operating System security patches.

Answer: D

Explanation:

In a VPC everything is allowed out by default.

QUESTION 193

Which of the following features ensures even distribution of traffic to Amazon EC2 instances in multiple Availability Zones registered with a load balancer?

- A. Elastic Load Balancing request routing
- B. An Amazon Route 53 weighted routing policy
- C. Elastic Load Balancing cross-zone load balancing
- D. An Amazon Route 53 latency routing policy

Answer: C

Explanation:

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/enable-disable-crosszone-lb.html>

QUESTION 194

You are using an m1.small EC2 Instance with one 300 GB EBS volume to host a relational database. You determined that write throughput to the database needs to be increased. Which of the following approaches can help achieve this? Choose 2 answers.

- A. Use an array of EBS volumes.
- B. Enable Multi-AZ mode.
- C. Place the instance in an Auto Scaling Groups
- D. Add an EBS volume and place into RAID 5.
- E. Increase the size of the EC2 Instance.
- F. Put the database behind an Elastic Load Balancer.

Answer: DE

QUESTION 195

After launching an instance that you intend to serve as a NAT (Network Address Translation) device in a public subnet you modify your route tables to have the NAT device be the target of internet bound traffic of your private subnet. When you try and make an outbound connection to the internet from an instance in the private subnet, you are not successful. Which of the following steps could resolve the issue?

- A. Disabling the Source/Destination Check attribute on the NAT instance
- B. Attaching an Elastic IP address to the instance in the private subnet
- C. Attaching a second Elastic Network Interface (ENI) to the NAT instance, and placing it in the private subnet
- D. Attaching a second Elastic Network Interface (ENI) to the instance in the private subnet, and placing it in the public subnet

Answer: A

Explanation:

http://docs.aws.amazon.com/workspaces/latest/adminguide/gsg_create_vpc.html

QUESTION 196

You are building a solution for a customer to extend their on-premises data center to AWS. The customer requires a 50-Mbps dedicated and private connection to their VPC. Which AWS product or feature satisfies this requirement?

- A. Amazon VPC peering
- B. Elastic IP Addresses
- C. AWS Direct Connect
- D. Amazon VPC virtual private gateway

Answer: C

QUESTION 197

You have multiple Amazon EC2 instances running in a cluster across multiple Availability Zones within the same region. What combination of the following should be used to ensure the highest network performance (packets per second), lowest latency, and lowest jitter? Choose 3 answers.

- A. Amazon EC2 placement groups
- B. Enhanced networking
- C. Amazon PV AMI
- D. Amazon HVM AMI
- E. Amazon Linux
- F. Amazon VPC

Answer: BDF

Explanation:

Enhanced Networking enables you to get significantly higher packet per second (PPS) performance, lower network jitter and lower latencies. This feature uses a new network virtualization stack that provides higher I/O performance and lower CPU utilization compared to traditional implementations. In order to take advantage of Enhanced Networking, you should launch an HVM AMI in VPC, and install the appropriate driver. For instructions on how to enable Enhanced Networking on EC2 instances, see the Enhanced Networking on Linux and Enhanced Networking on Windows tutorials. For availability of this feature by instance, or to learn more, visit the Enhanced Networking FAQ section.

QUESTION 198

When using the following AWS services, which should be implemented in multiple Availability Zones for high availability solutions? Choose 2 answers.

- A. Amazon DynamoDB
- B. Amazon Elastic Compute Cloud (EC2)
- C. Amazon Elastic Load Balancing
- D. Amazon Simple Notification Service (SNS)
- E. Amazon Simple Storage Service (S3)

Answer: BC

QUESTION 199

You have a video transcoding application running on Amazon EC2. Each instance polls a queue to find out which video should be transcoded, and then runs a transcoding process. If this process is interrupted, the video will be transcoded by another instance based on the queuing system. You have a large backlog of videos which need to be transcoded and would like to reduce this backlog by adding more instances. You will need these instances only until the backlog is reduced. Which type of Amazon EC2 instances should you use to reduce the backlog in the most cost efficient way?

- A. Reserved instances
- B. Spot instances
- C. Dedicated instances
- D. On-demand instances

Answer: B

Explanation:

<http://aws.amazon.com/ec2/purchasing-options/spot-instances/>

QUESTION 200

You have an EC2 Security Group with several running EC2 instances. You change the Security Group rules to allow inbound traffic on a new port and protocol, and launch several new instances in the same Security Group. The new rules apply:

- A. Immediately to all instances in the security group.
- B. Immediately to the new instances only.
- C. Immediately to the new instances, but old instances must be stopped and restarted before the new rules apply.
- D. To all instances, but it may take several minutes for old instances to see the changes.

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-network-security.html#vpc-security-groups>

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 201 – Question 250**

QUESTION 201

Which services allow the customer to retain full administrative privileges of the underlying EC2 instances? Choose 2 answers

- A. Amazon Relational Database Service
- B. Amazon Elastic Map Reduce
- C. Amazon ElastiCache
- D. Amazon DynamoDB
- E. AWS Elastic Beanstalk

Answer: BE

QUESTION 202

A company is building a two-tier web application to serve dynamic transaction-based content. The data tier is leveraging an Online Transactional Processing (OLTP) database. What services should you leverage to enable an elastic and scalable web tier?

- A. Elastic Load Balancing, Amazon EC2, and Auto Scaling
- B. Elastic Load Balancing, Amazon RDS with Multi-AZ, and Amazon S3
- C. Amazon RDS with Multi-AZ and Auto Scaling
- D. Amazon EC2, Amazon DynamoDB, and Amazon S3

Answer: A

QUESTION 203

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

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

Answer: C

QUESTION 204

Which technique can be used to integrate AWS IAM (Identity and Access Management) with an on-premise LDAP (Lightweight Directory Access Protocol) directory service?

- A. Use an IAM policy that references the LDAP account identifiers and the AWS credentials.
- B. Use SAML (Security Assertion Markup Language) to enable single sign-on between AWS and LDAP.
- C. Use AWS Security Token Service from an identity broker to issue short-lived AWS credentials.
- D. Use IAM roles to automatically rotate the IAM credentials when LDAP credentials are updated.
- E. Use the LDAP credentials to restrict a group of users from launching specific EC2 instance types.

Answer: B

Explanation:

<https://d0.awsstatic.com/whitepapers/aws-whitepaper-single-sign-on-integrating-aws-open-ldap-and-shibboleth.pdf>

QUESTION 205

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 IP.

Answer: BD

Explanation:

Even though we know the right Answers it is sometimes good to know why the other Answers are wrong.

- A. Is wrong because a subnet maps to a single AZ.
- C. Is wrong because /28 is the smallest subnet, amazon takes first four and last addresses per subnet.
- E. Is wrong because a private subnet needs a NAT appliance.

QUESTION 206

A customer is leveraging Amazon Simple Storage Service in eu-west-1 to store static content for a web-based property. The customer is storing objects using the Standard Storage class. Where are the customer's objects replicated?

- A. A single facility in eu-west-1 and a single facility in eu-central-1
- B. A single facility in eu-west-1 and a single facility in us-east-1
- C. Multiple facilities in eu-west-1
- D. A single facility in eu-west-1

Answer: C

Explanation:

Objects stored in a region never leave the region unless you explicitly transfer them to another region. For example, objects stored in the EU (Ireland) region never leave it.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/Introduction.html#Regions>

QUESTION 207

Your web application front end consists of multiple EC2 instances behind an Elastic Load Balancer. You configured ELB to perform health checks on these EC2 instances, if an instance fails to pass health checks, which statement will be true?

- A. The instance gets terminated automatically by the ELB.
- B. The instance gets quarantined by the ELB for root cause analysis.
- C. The instance is replaced automatically by the ELB.
- D. The ELB stops sending traffic to the instance that failed its health check.

Answer: D

QUESTION 208

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

Explanation:

http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

QUESTION 209

You have a web application running on six Amazon EC2 instances, consuming about 45% of resources on each instance. You are using auto-scaling to make sure that six instances are running at all times. The number of requests this application processes is consistent and does not experience spikes. The application is critical to your business and you want high availability at all times. You want the load to be distributed evenly between all instances. You also want to use the same Amazon Machine Image (AMI) for all instances. Which of the following architectural choices should you make?

- A. Deploy 6 EC2 instances in one availability zone and use Amazon Elastic Load Balancer.
- B. Deploy 3 EC2 instances in one region and 3 in another region and use Amazon Elastic Load Balancer.
- C. Deploy 3 EC2 instances in one availability zone and 3 in another availability zone and use Amazon Elastic Load Balancer.
- D. Deploy 2 EC2 instances in three regions and use Amazon Elastic Load Balancer.

Answer: C

Explanation:

A load balancer accepts incoming traffic from clients and routes requests to its registered EC2 instances in one or more Availability Zones.

<http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/how-elb-works.html>

Updated Security Whitepaper link:

<https://d0.awsstatic.com/whitepapers/aws-security-whitepaper.pdf>

QUESTION 210

You have decided to change the instance type for instances running in your application tier that is using Auto Scaling. In which area below would you change the instance type definition?

- A. Auto Scaling policy
- B. Auto Scaling group
- C. Auto Scaling tags
- D. Auto Scaling launch configuration

Answer: D

QUESTION 211

When an EC2 EBS-backed (EBS root) instance is stopped, what happens to the data on any ephemeral store volumes?

- A. Data is automatically saved in an EBS volume.
- B. Data is unavailable until the instance is restarted.
- C. Data will be deleted and will no longer be accessible.
- D. Data is automatically saved as an EBS snapshot.

Answer: C

Explanation:

An "EBS-backed" instance is an EC2 instance which uses an EBS volume as its root device.

An EBS volume behaves like a raw, unformatted, external block device that you can attach to a single instance and are not physically attached to the Instance host computer (more like a network attached storage). The volume persists independently from the running life of an instance. After an EBS volume is attached to an instance, you can use it like any other physical hard drive. You can also detach an EBS volume from one instance and attach it to another instance. EBS volumes can also be created as encrypted volumes using the Amazon EBS encryption feature.

QUESTION 212

Which of the following items are required to allow an application deployed on an EC2 instance to write data to a DynamoDB table? Assume that no security keys are allowed to be stored on the EC2 instance. Choose 2 answers.

- A. Create an IAM Role that allows write access to the DynamoDB table.
- B. Add an IAM Role to a running EC2 instance.
- C. Create an IAM User that allows write access to the DynamoDB table.
- D. Add an IAM User to a running EC2 instance.
- E. Launch an EC2 Instance with the IAM Role included in the launch configuration.

Answer: AE

Explanation:

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TicTacToe.Phase3.html>

QUESTION 213

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% durability. Therefore there is no need to confirm that data was inserted.
- C. A success code is inserted into the S3 object metadata.
- D. Each S3 account has a special bucket named `_s3_logs`. Success codes are written to this bucket with a timestamp and checksum.

Answer: A

Explanation:

To ensure that data is not corrupted traversing the network, use the Content-MD5 form field. When you use this form field, Amazon S3 checks the object against the provided MD5 value. If they do not match, Amazon S3 returns an error. The status code returned to the client upon successful upload if `success_action_redirect` is not specified. Accepts the values 200, 201, or 204 (default).

<http://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectPOST.html>

QUESTION 214

What is one key difference between an Amazon EBS-backed and an instance-store backed instance?

- A. Amazon EBS-backed instances can be stopped and restarted.
- B. Instance-store backed instances can be stopped and restarted.
- C. Auto scaling requires using Amazon EBS-backed instances.
- D. Virtual Private Cloud requires EBS backed instances.

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ComponentsAMIs.html#storage-for-the-root-device>

QUESTION 215

A company wants to implement their website in a virtual private cloud (VPC). The web tier will use an Auto Scaling group across multiple Availability Zones (AZs). The database will use Multi-AZ RDS MySQL and should not be publicly accessible. 'What is the minimum number of subnets that need to be configured in the VPC?

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

Answer: D

Explanation:

Would use VPC with private (DB) and public (WEB) subnets:

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.Scenarios.html

Multi AZ requirement forces me to multiply subnets by two.

Reasons:

For DB: Your VPC must have at least one subnet in at least two of the Availability Zones in the region where you want to deploy your DB instance. A subnet is a segment of a VPC's IP address range that you can specify and that lets you group instances based on your security and operational needs.

http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_VPC.WorkingWithRDSInstanceinaVPC.html

For Web: After creating a VPC, you can add one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span zones.

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 216

You have launched an Amazon Elastic Compute Cloud (EC2) instance into a public subnet with a primary private IP address assigned, an internet gateway is attached to the VPC, and the public route table is configured to send all Internet-based traffic to the Internet gateway. The instance security group is set to allow all outbound traffic but cannot access the internet. Why is the Internet unreachable from this instance?

- A. The instance does not have a public IP address.
- B. The internet gateway security group must allow all outbound traffic.
- C. The instance security group must allow all inbound traffic.
- D. The instance "Source/Destination check" property must be enabled.

Answer: A

Explanation:

Ensure that instances in your subnet have public IP addresses or Elastic IP addresses.

https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Internet_Gateway.html

QUESTION 217

You launch an Amazon EC2 instance without an assigned AVVS identity and Access Management (IAM) role. Later, you decide that the instance should be running with an IAM role. Which action must you take in order to have a running Amazon EC2 instance with an IAM role assigned to it?

- A. Create an image of the instance, and register the image with an IAM role assigned and an Amazon EBS volume mapping.
- B. Create a new IAM role with the same permissions as an existing IAM role, and assign it to the running instance.
- C. Create an image of the instance, add a new IAM role with the same permissions as the desired IAM role, and deregister the image with the new role assigned.
- D. Create an image of the instance, and use this image to launch a new instance with the desired IAM role assigned.

Answer: D

Explanation:

<http://docs.aws.amazon.com/IAM/latest/UserGuide/roles-usingrole-ec2instance.html>

QUESTION 218

How can the domain's zone apex, for example, "myzoneapexdomain.com", be pointed towards an Elastic Load Balancer?

- A. By using an Amazon Route 53 Alias record
- B. By using an AAAA record
- C. By using an Amazon Route 53 CNAME record
- D. By using an A record

Answer: A

Explanation:

You can create an alias resource record set at the zone apex. You cannot create a CNAME record at the top node of a DNS namespace, also known as the zone apex. For example, if you register the DNS name example.com, the zone apex is example.com.

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

QUESTION 219

An instance is launched into a VPC subnet with the network ACL configured to allow all inbound traffic and deny all outbound traffic. The instance's security group is configured to allow SSH from any IP address and deny all outbound traffic. What changes need to be made to allow SSH access to the instance?

- A. The outbound security group needs to be modified to allow outbound traffic.
- B. The outbound network ACL needs to be modified to allow outbound traffic.
- C. Nothing, it can be accessed from any IP address using SSH.
- D. Both the outbound security group and outbound network ACL need to be modified to allow outbound traffic.

Answer: B

Explanation:

Need to open TCP Port 1024-65535 at Outbound Rules

"Allows outbound responses to the remote computer. Network ACLs are stateless, therefore this rule is required to allow response traffic for inbound requests."

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_ACLS.html

QUESTION 220

For which of the following use cases are Simple Workflow Service (SWF) and Amazon EC2 an appropriate solution? Choose 2 answers.

- A. Using as an endpoint to collect thousands of data points per hour from a distributed fleet of sensors
- B. Managing a multi-step and multi-decision checkout process of an e-commerce website
- C. Orchestrating the execution of distributed and auditable business processes
- D. Using as an SNS (Simple Notification Service) endpoint to trigger execution of video transcoding jobs
- E. Using as a distributed session store for your web application

Answer: BC

Explanation:

<https://aws.amazon.com/swf/faqs/>

QUESTION 221

A customer wants to leverage Amazon Simple Storage Service (S3) and Amazon Glacier as part of their backup and archive infrastructure. The customer plans to use third-party software to support this integration. Which approach will limit the access of the third party software to only the Amazon S3 bucket named "company-backup"?

- A. A custom bucket policy limited to the Amazon S3 API in the Amazon Glacier archive "company-backup"
- B. A custom bucket policy limited to the Amazon S3 API in "company-backup"
- C. A custom IAM user policy limited to the Amazon S3 API for the Amazon Glacier archive "company-backup"
- D. A custom IAM user policy limited to the Amazon S3 API in "company-backup"

Answer: D

Explanation:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/example-policies-s3.html>

QUESTION 222

A client application requires operating system privileges on a relational database server. What is an appropriate configuration for a highly available database architecture?

- A. A standalone Amazon EC2 instance
- B. Amazon RDS in a Multi-AZ configuration
- C. Amazon EC2 instances in a replication configuration utilizing a single Availability Zone
- D. Amazon EC2 instances in a replication configuration utilizing two different Availability Zones

Answer: D

Explanation:

http://docs.aws.amazon.com/dms/latest/userguide/CHAP_Introduction.ReplicationInstance.html

QUESTION 223

What is a placement group?

- A. A collection of Auto Scaling groups in the same region
- B. A feature that enables EC2 instances to interact with each other via high bandwidth, low latency connections
- C. A collection of authorized CloudFront edge locations for a distribution
- D. A collection of Elastic Load Balancers in the same Region or Availability Zone

Answer: B

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gigabits per second (Gbps) network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

QUESTION 224

A company has a workflow that sends video files from their on-premise system to AWS for transcoding. They use EC2 worker instances that pull transcoding jobs from SQS. Why is SQS an appropriate service for this scenario?

- A. SQS guarantees the order of the messages.
- B. SQS synchronously provides transcoding output.
- C. SQS checks the health of the worker instances.
- D. SQS helps to facilitate horizontal scaling of encoding tasks.

Answer: D

Explanation:

Imho the idea for SQS is to improve scalability.

Elastic Beanstalk is checking the health of EC2 instances, not sure if SQS does.

QUESTION 225

When creation of an EBS snapshot is initiated, but not completed, the EBS volume:

- A. Can be used while the snapshot is in progress.
- B. Cannot be detached or attached to an EC2 instance until the snapshot completes.
- C. Can be used in read-only mode while the snapshot is in progress.
- D. Cannot be used until the snapshot completes.

Answer: A

Explanation:

Snapshots occur asynchronously; the point-in-time snapshot is created immediately, but the status of the snapshot is pending until the snapshot is complete (when all of the modified blocks have been transferred to Amazon S3), which can take several hours for large initial snapshots or subsequent snapshots where many blocks have changed.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-creating-snapshot.html>

QUESTION 226

What are characteristics of Amazon S3? Choose 2 answers.

- A. S3 allows you to store objects of virtually unlimited size.
- B. S3 offers Provisioned IOPS.
- C. S3 allows you to store unlimited amounts of data.
- D. S3 should be used to host a relational database.
- E. Objects are directly accessible via a URL.

Answer: CE

Explanation:

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>

QUESTION 227

Per the AWS Acceptable Use Policy, penetration testing of EC2 instances:

- A. May be performed by AWS, and will be performed by AWS upon customer request.
- B. May be performed by AWS, and is periodically performed by AWS.
- C. Are expressly prohibited under all circumstances.
- D. May be performed by the customer on their own instances with prior authorization from AWS.
- E. May be performed by the customer on their own instances, only if performed from EC2 instances.

Answer: D

Explanation:

<http://aws.amazon.com/security/penetration-testing/>

QUESTION 228

You are working with a customer who has 10 TB of archival data that they want to migrate to Amazon Glacier. The customer has a 1-Mbps connection to the Internet. Which service or feature provides the fastest method of getting the data into Amazon Glacier?

- A. Amazon Glacier multipart upload
- B. AWS Storage Gateway
- C. VM Import/Export
- D. AWS Import/Export

Answer: A

Explanation:

You can only perform an Amazon Glacier import from devices of 4 TB in size or smaller.

https://docs.aws.amazon.com/es_es/AWSImportExport/latest/DG/createGlacierimportjobs.html

QUESTION 229

How can you secure data at rest on an EBS volume?

- A. Attach the volume to an instance using EC2's SSL interface.
- B. Write the data randomly instead of sequentially.
- C. Encrypt the volume using the S3 server-side encryption service.
- D. Create an IAM policy that restricts read and write access to the volume.
- E. Use an encrypted file system on top of the EBS volume.

Answer: E

Explanation:

<https://aws.amazon.com/blogs/aws/protect-your-data-with-new-ebs-encryption/>

QUESTION 230

A customer needs to capture all client connection information from their load balancer every five minutes. The company wants to use this data for analyzing traffic patterns and troubleshooting their applications. Which of the following options meets the customer requirements?

- A. Enable AWS CloudTrail for the load balancer.
- B. Enable access logs on the load balancer.
- C. Install the Amazon CloudWatch Logs agent on the load balancer.
- D. Enable Amazon CloudWatch metrics on the load balancer.

Answer: B

Explanation:

Elastic Load Balancing access logs

The access logs for Elastic Load Balancing capture detailed information for all requests made to your load balancer and stores them as log files in the Amazon S3 bucket that you specify. Each log contains details such as the time a request was received, the client's IP address, latencies, request path, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot your back-end applications. For more information, see [Monitor Your Load Balancer Using Elastic Load Balancing Access Logs](#).

QUESTION 231

If you want to launch Amazon Elastic Compute Cloud (EC2) instances and assign each instance a predetermined private IP address you should:

- A. Launch the instance from a private Amazon Machine Image (AMI).
- B. Assign a group of sequential Elastic IP address to the instances.
- C. Launch the instances in the Amazon Virtual Private Cloud (VPC).
- D. Launch the instances in a Placement Group.
- E. Use standard EC2 instances since each instance gets a private Domain Name Service (DNS) already.

Answer: C

Explanation:

Each instance in a VPC has a default network interface (eth0) that is assigned the primary private IP address.

QUESTION 232

You need to configure an Amazon S3 bucket to serve static assets for your public-facing web application. Which methods ensure that all objects uploaded to the bucket are set to public read? Choose 2 answers.

- A. Set permissions on the object to public read during upload.
- B. Configure the bucket ACL to set all objects to public read.
- C. Configure the bucket policy to set all objects to public read.
- D. Use AWS Identity and Access Management roles to set the bucket to public read.
- E. Amazon S3 objects default to public read, so no action is needed.

Answer: AC

Explanation:

<https://aws.amazon.com/articles/5050>

You can use ACLs to grant permissions to individual AWS accounts; however, it is strongly recommended that you do not grant public access to your bucket using an ACL. So the recommended approach is create bucket policy, but not ACL. Following link give you an example about how to make the bucket content public.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/HostingWebsiteOnS3Setup.html#step2-add-bucket-policy-make-content-public>

QUESTION 233

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

Explanation:

<http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingKMSEncryption.html>

QUESTION 234

Which procedure for backing up a relational database on EC2 that is using a set of RAIDed EBS volumes for storage minimizes the time during which the database cannot be written to and results in a consistent backup?

- A. 1. Detach EBS volumes, 2. Start EBS snapshot of volumes, 3. Re-attach EBS volumes
- B. 1. Stop the EC2 Instance, 2. Snapshot the EBS volumes
- C. 1. Suspend disk I/O, 2. Create an image of the EC2 Instance, 3. Resume disk I/O
- D. 1. Suspend disk I/O, 2. Start EBS snapshot of volumes, 3. Resume disk I/O
- E. 1. Suspend disk I/O, 2. Start EBS snapshot of volumes, 3. Wait for snapshots to complete, 4. Resume disk I/O

Answer: B

Explanation:

<https://aws.amazon.com/cn/premiumsupport/knowledge-center/snapshot-ebs-raid-array/>

To create an "application-consistent" snapshot of your RAID array, stop applications from writing to the RAID array, and flush all caches to disk. Then ensure that the associated EC2 instance is no longer writing to the RAID array by taking steps such as freezing the file system, unmounting the RAID array, or *shutting down the associated EC2 instance*. After completing the steps to halt all I/O, take a snapshot of each EBS volume.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-detaching-volume.html>

You can detach an Amazon EBS volume from an instance explicitly or by terminating the instance. However, if the instance is running, you must first unmount the volume from the instance."

QUESTION 235

A company needs to deploy virtual desktops to its customers in a virtual private cloud, leveraging existing security controls. Which set of AWS services and features will meet the company's requirements?

- A. Virtual Private Network connection, AWS Directory Services, and ClassicLink
- B. Virtual Private Network connection, AWS Directory Services, and Amazon Workspaces
- C. AWS Directory Service, Amazon Workspaces, and AWS Identity and Access Management
- D. Amazon Elastic Compute Cloud, and AWS Identity and Access Management

Answer: B

Explanation:

To enable integration, you need to ensure that your domain is reachable via an Amazon Virtual Private Cloud VPC (this could mean that Active Directory domain controllers for your domain are running on Amazon EC2 instances, or that they are reachable via a VPN connection and are located in your on-premises network).

QUESTION 236

After creating a new IAM user which of the following must be done before they can successfully make API calls?

- A. Add a password to the user.
- B. Enable Multi-Factor Authentication for the user.
- C. Assign a Password Policy to the user.
- D. Create a set of Access Keys for the user.

Answer: D

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_SettingUpUser.html

QUESTION 237

Which of the following are valid statements about Amazon S3? Choose 2 answers.

- A. S3 provides read-after-write consistency for any type of PUT or DELETE.
- B. Consistency is not guaranteed for any type of PUT or DELETE.
- C. A successful response to a PUT request only occurs when a complete object is saved.
- D. Partially saved objects are immediately readable with a GET after an overwrite PUT.
- E. S3 provides eventual consistency for overwrite PUTS and DELETES.

Answer: CE

Explanation:

<http://api-portal.anypoint.mulesoft.com/amazon/api/amazon-s3-api/docs/concepts#DataConsistencyModel>

QUESTION 238

You are configuring your company's application to use Auto Scaling and need to move user state information. Which of the following AWS services provides a shared data store with durability and low latency?

- A. AWS ElastiCache Memcached
- B. Amazon Simple Storage Service
- C. Amazon EC2 instance storage
- D. Amazon DynamoDB

Answer: D

Explanation:

https://media.amazonwebservices.com/AWS_Storage_Options.pdf

To speed access to relevant data, many developers pair Amazon S3 with a database, such as Amazon DynamoDB or Amazon RDS. Amazon S3 stores the actual information, and the database serves as the repository for associated metadata (e.g., object name, size, keywords, and so on). Metadata in the database can easily be indexed and queried, making it very efficient to locate an object's reference via a database query. This result can then be used to pinpoint and then retrieve the object itself from Amazon S3.

QUESTION 239

Which features can be used to restrict access to data in S3? Choose 2 answers.

- A. Set an S3 ACL on the bucket or the object.
- B. Create a CloudFront distribution for the bucket.
- C. Set an S3 bucket policy.
- D. Enable IAM Identity Federation.
- E. Use S3 Virtual Hosting.

Answer: AC

Explanation:

Amazon S3 is secure by default. Only the bucket and object owners originally have access to Amazon S3 resources they create. Amazon S3 supports user authentication to control access to data. You can use access control mechanisms such as bucket policies and Access Control Lists (ACLs) to selectively grant permissions to users and groups of users. You can securely upload/download your data to Amazon S3 via SSL endpoints using the HTTPS protocol. If you need extra security you can use the Server Side Encryption (SSE) option or the Server Side Encryption with Customer-Provide Keys (SSE-C) option to encrypt data stored-at-rest. Amazon S3 provides the encryption technology for both SSE and SSE-C. Alternatively you can use your own encryption libraries to encrypt data before storing it in Amazon S3.

<https://aws.amazon.com/s3/faqs/>

QUESTION 240

Which of the following are characteristics of a reserved instance? Choose 3 answers.

- A. It can be migrated across Availability Zones
- B. It is specific to an Amazon Machine Image (AMI)
- C. It can be applied to instances launched by Auto Scaling
- D. It is specific to an instance Type

E. It can be used to lower Total Cost of Ownership (TCO) of a system

Answer: ACE

Explanation:

You can use Auto Scaling or other AWS services to launch the On-Demand instances that use your Reserved Instance benefits. For information about launching On-Demand instances, see Launch Your Instance. For information about launching instances using Auto Scaling, see the Auto Scaling User Guide.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts-on-demand-reserved-instances.html>

<https://forums.aws.amazon.com/thread.jspa?threadID=56501>

QUESTION 241

Which Amazon Elastic Compute Cloud feature can you query from within the instance to access instance properties?

- A. Instance user data
- B. Resource tags
- C. Instance metadata
- D. Amazon Machine Image

Answer: C

Explanation:

Although you can only access instance metadata and user data from within the instance itself, the data is not protected by cryptographic methods.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-metadata.html#instancedata-data-retrieval>

QUESTION 242

Which of the following requires a custom CloudWatch metric to monitor?

- A. Memory Utilization of an EC2 instance
- B. CPU Utilization of an EC2 instance
- C. Disk usage activity of an EC2 instance
- D. Data transfer of an EC2 instance

Answer: A

Explanation:

CloudWatch relies on the information provided by this hypervisor, which can only see the most hardware-sided part of the instance's status, including CPU usage (but not load), total memory size (but not memory usage), number of I/O operations on the hard disks (but not its partition layout and space usage) and network traffic (but not the processes generating it).

QUESTION 243

You are tasked with setting up a Linux bastion host for access to Amazon EC2 instances running in your VPC. Only clients connecting from the corporate external public IP address 72.34.51.100 should have SSH access to the host. Which option will meet the customer requirement?

- A. Security Group Inbound Rule: Protocol -TCP, Port Range -22, Source 72.34.51.100/32
- B. Security Group Inbound Rule: Protocol -UDP, Port Range -22, Source 72.34.51.100/32
- C. Network ACL Inbound Rule: Protocol -UDP, Port Range -22, Source 72.34.51.100/32
- D. Network ACL Inbound Rule: Protocol -TCP, Port Range-22, Source 72.34.51.100/0

Answer: A

QUESTION 244

A customer needs corporate IT governance and cost oversight of all AWS resources consumed by its divisions. The divisions want to maintain administrative control of the discrete AWS resources they consume and keep those resources separate from the resources of other divisions. Which of the following options, when used together will support the autonomy/control of divisions while enabling corporate IT to maintain governance and cost oversight? Choose 2 answers.

- A. Use AWS Consolidated Billing and disable AWS root account access for the child accounts.
- B. Enable IAM cross-account access for all corporate IT administrators in each child account.
- C. Create separate VPCs for each division within the corporate IT AWS account.
- D. Use AWS Consolidated Billing to link the divisions' accounts to a parent corporate account.
- E. Write all child AWS CloudTrail and Amazon CloudWatch logs to each child account's Amazon S3 'Log' bucket.

Answer: BD

Explanation:

http://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

<http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html>

QUESTION 245

You run an ad-supported photo sharing website using S3 to serve photos to visitors of your site. At some point you find out that other sites have been linking to the photos on your site, causing loss to your business. What is an effective method to mitigate this?

- A. Remove public read access and use signed URLs with expiry dates.
- B. Use CloudFront distributions for static content.
- C. Block the IPs of the offending websites in Security Groups.
- D. Store photos on an EBS volume of the web server.

Answer: A

Explanation:

A signed URL includes additional information, for example, an expiration date and time, that gives you more control over access to your content.

QUESTION 246

You are working with a customer who is using Chef configuration management in their data center. Which service is designed to let the customer leverage existing Chef recipes in AWS?

- A. Amazon Simple Workflow Service
- B. AWS Elastic Beanstalk
- C. AWS CloudFormation
- D. AWS OpsWorks

Answer: D

Explanation:

<http://aws.amazon.com/opsworks/>

QUESTION 247

An Auto-Scaling group spans 3 AZs and currently has 4 running EC2 instances. When Auto Scaling needs to terminate an EC2 instance by default, AutoScaling will: _____. Choose 2 answers.

- A. Allow at least five minutes for Windows/Linux shutdown scripts to complete, before terminating the instance.
- B. Terminate the instance with the least active network connections. If multiple instances meet this criterion, one will be randomly selected.
- C. Send an SNS notification, if configured to do so.
- D. Terminate an instance in the AZ which currently has 2 running EC2 instances.
- E. Randomly select one of the 3 AZs, and then terminate an instance in that AZ.

Answer: CD

Explanation:

Auto Scaling determines whether there are instances in multiple Availability Zones. If so, it selects the Availability Zone with the most instances and at least one instance that is not protected from scale in.

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/AutoScalingBehavior.InstanceTermination.html>

QUESTION 248

When an EC2 instance that is backed by an S3-based AMI is terminated, what happens to the data on the root volume?

- A. Data is automatically saved as an EBS snapshot.
- B. Data is automatically saved as an EBS volume.
- C. Data is unavailable until the instance is restarted.
- D. Data is automatically deleted.

Answer: D

Explanation:

Using the legacy S3 based AMIs, either of the above terminates the instance and you lose all local and ephemeral storage (boot disk and /mnt) forever. Hope you remembered to save the important stuff elsewhere.

QUESTION 249

In order to optimize performance for a compute cluster that requires low inter-node latency, which of the following feature should you use?

- A. Multiple Availability Zones
- B. AWS Direct Connect
- C. EC2 Dedicated Instances
- D. Placement Groups
- E. VPC private subnets

Answer: D

Explanation:

A placement group is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to participate in a low-latency, 10 Gigabits per second (Gbps) network.
<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html>

QUESTION 250

You have an environment that consists of a public subnet using Amazon VPC and 3 instances that are running in this subnet. These three instances can successfully communicate with other hosts on the Internet. You launch a fourth instance in the same subnet, using the same AMI and security group configuration you used for the others, but find that this instance cannot be accessed from the internet. What should you do to enable Internet access?

- A. Deploy a NAT instance into the public subnet.
- B. Assign an Elastic IP address to the fourth instance.
- C. Configure a publically routable IP Address in the host OS of the fourth instance.
- D. Modify the routing table for the public subnet.

Answer: B

Explanation:

You launched your instance into a public subnet - a subnet that has a route to an Internet gateway. However, the instance in your subnet also needs a public IP address to be able to communicate with the Internet. By default, an instance in a nondefault VPC is not assigned a public IP address. In this step, you'll allocate an Elastic IP address to your account, and then associate it with your instance

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 251 – Question 300**

QUESTION 251

You have a distributed application that periodically processes large volumes of data across multiple Amazon EC2 Instances. The application is designed to recover gracefully from Amazon EC2 instance failures. You are required to accomplish this task in the most cost-effective way. Which of the following will meet your requirements?

- A. Spot Instances
- B. Reserved instances
- C. Dedicated instances
- D. On-Demand instances

Answer: A

Explanation:

Using reserved instances is not the most cost-effective way.

<https://aws.amazon.com/blogs/aws/new-scheduled-reserved-instances/>

“Scheduled Reserved Instance model allows you to reserve instances for predefined blocks of time on a recurring basis for a one-year term, with prices that are generally 5 to 10% lower than the equivalent On-Demand rates.”

You can get spot instances with much lower prices:

<https://aws.amazon.com/ec2/spot/pricing/>

“Spot instances are also available to run for a predefined duration – in hourly increments up to six hours in length – at a significant discount (30-45%) compared to On-Demand pricing plus an additional 5% during off-peak times for a total of up to 50% savings.”

QUESTION 252

Which of the following are true regarding AWS CloudTrail? Choose 3 answers.

- A. CloudTrail is enabled globally.
- B. CloudTrail is enabled by default.
- C. CloudTrail is enabled on a per-region basis.
- D. CloudTrail is enabled on a per-service basis.
- E. Logs can be delivered to a single Amazon S3 bucket for aggregation.
- F. CloudTrail is enabled for all available services within a region.
- G. Logs can only be processed and delivered to the region in which they are generated.

Answer: ACE

Explanation:

A: have a trail with the Apply trail to all regions option enabled.

C: have multiple single region trails.

E: Log files from all the regions can be delivered to a single S3 bucket.

Global service events are always delivered to trails that have the Apply trail to all regions option enabled. Events are delivered from a single region to the bucket for the trail. This setting cannot be changed.

If you have a single region trail, you should enable the Include global services option.

If you have multiple single region trails, you should enable the Include global services option in only one of the trails. D: Incorrect. Once enabled it is applicable for all the supported services, service can't be selected.

QUESTION 253

You have a content management system running on an Amazon EC2 instance that is approaching 100% CPU utilization. Which option will reduce load on the Amazon EC2 instance?

- A. Create a load balancer, and register the Amazon EC2 instance with it
- B. Create a CloudFront distribution, and configure the Amazon EC2 instance as the origin
- C. Create an Auto Scaling group from the instance using the CreateAutoScalingGroup action
- D. Create a launch configuration from the instance using the CreateLaunchConfiguration action

Answer: C

Explanation:

You can create an ASG from instance ID

http://docs.aws.amazon.com/AutoScaling/latest/APIReference/API_CreateAutoScalingGroup.html

QUESTION 254

You have a load balancer configured for VPC, and all back-end Amazon EC2 instances are in service. However, your web browser times out when connecting to the load balancer's DNS name. Which options are probable causes of this behavior? Choose 2 answers.

- A. The load balancer was not configured to use a public subnet with an Internet gateway configured.
- B. The Amazon EC2 instances do not have a dynamically allocated private IP address.
- C. The security groups or network ACLs are not properly configured for web traffic.
- D. The load balancer is not configured in a private subnet with a NAT instance.
- E. The VPC does not have a VGW configured.

Answer: AC

Explanation:

There is no such thing as VGW. Hence E is not correct answer.

QUESTION 255

A company needs to deploy services to an AWS region which they have not previously used. The company currently has an AWS identity and Access Management (IAM) role for the Amazon EC2 instances, which permits the instance to have access to Amazon DynamoDB. The company wants their EC2 instances in the new region to have the same privileges. How should the company achieve this?

- A. Create a new IAM role and associated policies within the new region
- B. Assign the existing IAM role to the Amazon EC2 instances in the new region
- C. Copy the IAM role and associated policies to the new region and attach it to the instances
- D. Create an Amazon Machine Image (AMI) of the instance and copy it to the desired region using the AMI Copy feature

Answer: B

QUESTION 256

Which of the following notification endpoints or clients are supported by Amazon Simple Notification Service? Choose 2 answers.

- A. Email
- B. CloudFront distribution
- C. File Transfer Protocol
- D. Short Message Service
- E. Simple Network Management Protocol

Answer: AD

Explanation:

SNS Supported Endpoints

Email Notifications

Amazon SNS provides the ability to send Email notifications.

SMS Notifications

Amazon SNS provides the ability to send and receive Short Message Service (SMS) notifications to SMS-enabled mobile phones and smart phones.

<http://docs.aws.amazon.com/sns/latest/dg/welcome.html>

QUESTION 257

Which set of Amazon S3 features helps to prevent and recover from accidental data loss?

- A. Object lifecycle and service access logging
- B. Object versioning and Multi-factor authentication
- C. Access controls and server-side encryption
- D. Website hosting and Amazon S3 policies

Answer: B

Explanation:

Versioning-enabled buckets enable you to recover objects from accidental deletion or overwrite. In addition to that, they have made it a requirement that delete operations on versioned data can only be done using MFA (Multi factor authentication).

http://media.amazonwebservices.com/AWS_Security_Best_Practices.pdf

QUESTION 258

A company needs to monitor the read and write IOPs metrics for their AWS MySQL RDS instance and send real-time alerts to their operations team. Which AWS services can accomplish this? Choose 2 answers.

- A. Amazon Simple Email Service
- B. Amazon CloudWatch
- C. Amazon Simple Queue Service
- D. Amazon Route 53
- E. Amazon Simple Notification Service

Answer: BE

Explanation:

B: Amazon RDS provides metrics in real time for the operating system (OS) that your DB instance runs on. You can view the metrics for your DB instance using the console, or consume the Enhanced Monitoring JSON output from CloudWatch Logs in a monitoring system of your choice.

E: Use Amazon RDS DB events to monitor failovers. For example, you can be notified by text message or email when a DB instance fails over. Amazon RDS uses the Amazon Simple Notification Service (Amazon SNS) to provide notification when an Amazon RDS event occurs.

QUESTION 259

A company is preparing to give AWS Management Console access to developers. Company policy mandates identity federation and role-based access control. Roles are currently assigned using groups in the corporate Active Directory. What combination of the following will give developers access to the AWS console? Choose 2 answers.

- A. AWS Directory Service AD Connector
- B. AWS Directory Service Simple AD
- C. AWS Identity and Access Management groups
- D. AWS identity and Access Management roles
- E. AWS identity and Access Management users

Answer: AD

Explanation:

http://docs.aws.amazon.com/STS/latest/APIReference/API_AssumeRoleWithSAML.html

QUESTION 260

What is the durability of S3 RRS?

- A. 99.99%
- B. 99.95%
- C. 99.995%
- D. 99.999999999%

Answer: A

QUESTION 261

What does specifying the mapping `/dev/sdc=none` when launching an instance do? [html](#)

- A. Prevents /dev/sdc from creating the instance.
- B. Prevents /dev/sdc from deleting the instance.
- C. Set the value of /dev/sdc to 'zero'.
- D. Prevents /dev/sdc from attaching to the instance.

Answer: D

QUESTION 262

You are deploying an application to track GPS coordinates of delivery trucks in the United States. Coordinates are transmitted from each delivery truck once every three seconds. You need to design an architecture that will enable real-time processing of these coordinates from multiple consumers. Which service should you use to implement data ingestion?

- A. Amazon Kinesis
- B. AWS Data Pipeline
- C. Amazon AppStream
- D. Amazon Simple Queue Service

Answer: A

Explanation:

<https://aws.amazon.com/streaming-data/>

QUESTION 263

A photo-sharing service stores pictures in Amazon Simple Storage Service (S3) and allows application sign-in using an OpenID Connect-compatible identity provider. Which AWS Security Token Service approach to temporary access should you use for the Amazon S3 operations?

- A. SAML-based Identity Federation
- B. Cross-Account Access
- C. AWS Identity and Access Management roles
- D. Web Identity Federation

Answer: D

Explanation:

Web identity federation - You can let users sign in using a well-known third party identity provider such as Login with Amazon, Facebook, Google, or any OpenID Connect (OIDC) 2.0 compatible provider. AWS STS web identity federation supports Login with Amazon, Facebook, Google, and any OpenID Connect (OIDC)-compatible identity provider.

QUESTION 264

You have an application running on an Amazon Elastic Compute Cloud instance, that uploads 5 GB video objects to Amazon Simple Storage Service (S3). Video uploads are taking longer than expected, resulting in poor application performance. Which method will help improve performance of your application?

- A. Enable enhanced networking
- B. Use Amazon S3 multipart upload
- C. Leveraging Amazon CloudFront, use the HTTP POST method to reduce latency
- D. Use Amazon Elastic Block Store Provisioned IOPs and use an Amazon EBS-optimized instance

Answer: B

Explanation:

Using multipart upload provides the following advantages:

- Improved throughput - You can upload parts in parallel to improve throughput.
- Quick recovery from any network issues - Smaller part size minimizes the impact of restarting a failed upload due to a network error.
- Pause and resume object uploads - You can upload object parts over time. Once you initiate a multipart upload there is no expiry; you must explicitly complete or abort the multipart upload.
- Begin an upload before you know the final object size.
- You can upload an object as you are creating it.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/uploadobjusingmpu.html>

QUESTION 265

A customer wants to track access to their Amazon Simple Storage Service (S3) buckets and also use this information for their internal security and access audits. Which of the following will meet the Customer requirement?

- A. Enable AWS CloudTrail to audit all Amazon S3 bucket access.
- B. Enable server access logging for all required Amazon S3 buckets.
- C. Enable the Requester Pays option to track access via AWS Billing.
- D. Enable Amazon S3 event notifications for Put and Post.

Answer: B

Explanation:

If it's just for internal audit, then Server access logging, I assume is sufficient.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/ServerLogs.html>

For external audits I would go for CloudTrail.

<http://docs.aws.amazon.com/AmazonS3/latest/dev/cloudtrail-logging.html>

QUESTION 266

A company is deploying a two-tier, highly available web application to AWS. Which service provides durable storage for static content while utilizing lower Overall CPU resources for the web tier?

- A. Amazon EBS volume
- B. Amazon S3
- C. Amazon EC2 instance store
- D. Amazon RDS instance

Answer: B

QUESTION 267

You are designing a web application that stores static assets in an Amazon Simple Storage Service (S3) bucket. You expect this bucket to immediately receive over 150 PUT requests per second. What should you do to ensure optimal performance?

- A. Use multi-part upload.
- B. Add a random prefix to the key names.
- C. Amazon S3 will automatically manage performance at this scale.
- D. Use a predictable naming scheme, such as sequential numbers or date time sequences, in the key names.

Answer: B

Explanation:

If you anticipate that your workload will consistently exceed 100 requests per second, you should avoid sequential key names. If you must use sequential numbers or date and time patterns in key names, add a random prefix to the key name. The randomness of the prefix more evenly distributes key names across multiple index partitions. Examples of introducing randomness are provided later in this topic.

QUESTION 268

When will you incur costs with an Elastic IP address (EIP)?

- A. When an EIP is allocated.
- B. When it is allocated and associated with a running instance.
- C. When it is allocated and associated with a stopped instance.
- D. Costs are incurred regardless of whether the EIP is associated with a running instance.

Answer: C

Explanation:

You are allowed one EIP to be attached to a running instance at no charge. otherwise, it will incur a small fee. in this case, the instance is stopped, and thus, the EIP will be billed at the normal rate.

<http://aws.amazon.com/ec2/pricing/>

QUESTION 269

A company has an AWS account that contains three VPCs (Dev, Test, and Prod) in the same region. Test is peered to both Prod and Dev. All VPCs have non-overlapping CIDR blocks. The company wants to push minor code releases from Dev to Prod to speed up time to market. Which of the following options helps the company accomplish this?

- A. Create a new peering connection Between Prod and Dev along with appropriate routes.

- B. Create a new entry to Prod in the Dev route table using the peering connection as the target.
- C. Attach a second gateway to Dev. Add a new entry in the Prod route table identifying the gateway as the target.
- D. The VPCs have non-overlapping CIDR blocks in the same account. The route tables contain local routes for all VPCs.

Answer: A

Explanation:

<http://docs.aws.amazon.com/AmazonVPC/latest/PeeringGuide/vpc-pg.pdf#create-vpc-peering-connection>

QUESTION 270

Which of the following instance types are available as Amazon EBS-backed only? Choose 2 answers.

- A. General purpose T2
- B. General purpose M3
- C. Compute-optimized C4
- D. Compute-optimized C3
- E. Storage-optimized 12

Answer: AC

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html>

QUESTION 271

A customer is hosting their company website on a cluster of web servers that are behind a public-facing load balancer. The customer also uses Amazon Route 53 to manage their public DNS. How should the customer configure the DNS zone apex record to point to the load balancer?

- A. Create an A record pointing to the IP address of the load balancer.
- B. Create a CNAME record pointing to the load balancer DNS name.
- C. Create a CNAME record aliased to the load balancer DNS name.
- D. Create an A record aliased to the load balancer DNS name.

Answer: D

QUESTION 272

You try to connect via SSH to a newly created Amazon EC2 instance and get one of the following error messages: "Network error: Connection timed out" or "Error connecting to [instance], reason: -> Connection timed out: connect," You have confirmed that the network and security group rules are configured correctly and the instance is passing status checks. What steps should you take to identify the source of the behavior? Choose 2 answers.

- A. Verify that the private key file corresponds to the Amazon EC2 key pair assigned at launch.
- B. Verify that your IAM user policy has permission to launch Amazon EC2 instances.
- C. Verify that you are connecting with the appropriate user name for your AMI.
- D. Verify that the Amazon EC2 Instance was launched with the proper IAM role.
- E. Verify that your federation trust to AWS has been established.

Answer: AC

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstancesConnecting.html>

QUESTION 273

A customer is running a multi-tier web application farm in a virtual private cloud (VPC) that is not connected to their corporate network. They are connecting to the VPC over the Internet to manage all of their Amazon EC2 instances running in both the public and private subnets. They have only authorized the bastion-security-group with Microsoft Remote Desktop Protocol (RDP) access to the application instance security groups, but the company wants to further limit administrative access to all of the instances in the VPC. Which of the following Bastion deployment scenarios will meet this requirement?

- A. Deploy a Windows Bastion host on the corporate network that has RDP access to all instances in the VPC.
- B. Deploy a Windows Bastion host with an Elastic IP address in the public subnet and allow SSH access to the bastion from anywhere.
- C. Deploy a Windows Bastion host with an Elastic IP address in the private subnet, and restrict RDP access to the bastion from only the corporate public IP addresses.

- D. Deploy a Windows Bastion host with an auto-assigned Public IP address in the public subnet, and allow RDP access to the bastion from only the corporate public IP addresses.

Answer: D

QUESTION 274

A customer has a single 3-TB volume on-premises that is used to hold a large repository of images and print layout files. This repository is growing at 500 GB a year and must be presented as a single logical volume. The customer is becoming increasingly constrained with their local storage capacity and wants an off-site backup of this data, while maintaining low-latency access to their frequently accessed data. Which AWS Storage Gateway configuration meets the customer requirements?

- A. Gateway-Cached volumes with snapshots scheduled to Amazon S3
- B. Gateway-Stored volumes with snapshots scheduled to Amazon S3
- C. Gateway-Virtual Tape Library with snapshots to Amazon S3
- D. Gateway-Virtual Tape Library with snapshots to Amazon Glacier

Answer: A

Explanation:

<http://docs.aws.amazon.com/storagegateway/latest/userguide/storage-gateway-cached-concepts.html>

QUESTION 275

You are building an automated transcription service in which Amazon EC2 worker instances process an uploaded audio file and generate a text file. You must store both of these files in the same durable storage until the text file is retrieved. You do not know what the storage capacity requirements are. Which storage option is both cost-efficient and scalable?

- A. Multiple Amazon EBS volume with snapshots
- B. A single Amazon Glacier vault
- C. A single Amazon S3 bucket
- D. Multiple instance stores

Answer: C

QUESTION 276

You need to pass a custom script to new Amazon Linux instances created in your Auto Scaling group. Which feature allows you to accomplish this?

- A. User data
- B. EC2Config service
- C. IAM roles
- D. AWS Config

Answer: A

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/user-data.html#user-data-shell-scripts>

Not B, because EC2Config is used for Windows instances:

http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/UsingConfig_WinAMI.html

QUESTION 277

Which of the following services natively encrypts data at rest within an AWS region? Choose 2 answers.

- A. AWS Storage Gateway
- B. Amazon DynamoDB
- C. Amazon CloudFront
- D. Amazon Glacier
- E. Amazon Simple Queue Service

Answer: AD

Explanation:

https://media.amazonwebservices.com/AWS_Securing_Data_at_Rest_with_Encryption.pdf (page 12)

QUESTION 278

A company is building software on AWS that requires access to various AWS services. Which configuration should be used to ensure mat AWS credentials (i.e., Access Key ID/Secret Access Key combination) are not compromised?

- A. Enable Multi-Factor Authentication for your AWS root account.
- B. Assign an IAM role to the Amazon EC2 instance.
- C. Store the AWS Access Key ID/Secret Access Key combination in software comments.
- D. Assign an IAM user to the Amazon EC2 Instance.

Answer: B

Explanation:

Use roles for applications that run on Amazon EC2 instances.

Applications that run on an Amazon EC2 instance need credentials in order to access other AWS services. To provide credentials to the application in a secure way, use IAM roles. A role is an entity that has its own set of permissions, but that isn't a user or group. Roles also don't have their own permanent set of credentials the way IAM users do. In the case of Amazon EC2, IAM dynamically provides temporary credentials to the EC2 instance, and these credentials are automatically rotated for you.

<http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#use-roles-with-ec2>

QUESTION 279

Which of the following are true regarding encrypted Amazon Elastic Block Store (EBS) volumes? Choose 2 answers.

- A. Supported on all Amazon EBS volume types
- B. Snapshots are automatically encrypted
- C. Available to all instance types
- D. Existing volumes can be encrypted
- E. shared volumes can be encrypted

Answer: AB

Explanation:

This feature is supported on all Amazon EBS volume types (General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic). You can access encrypted Amazon EBS volumes the same way you access existing volumes; encryption and decryption are handled transparently and they require no additional action from you, your Amazon EC2 instance, or your application. Snapshots of encrypted Amazon EBS volumes are automatically encrypted, and volumes that are created from encrypted Amazon EBS snapshots are also automatically encrypted.

<http://docs.aws.amazon.com/kms/latest/developerguide/services-ebs.html>

QUESTION 280

A company is deploying a new two-tier web application in AWS. The company has limited staff and requires high availability, and the application requires complex queries and table joins. Which configuration provides the solution for the company's requirements?

- A. MySQL Installed on two Amazon EC2 Instances in a single Availability Zone
- B. Amazon RDS for MySQL with Multi-AZ
- C. Amazon ElastiCache
- D. Amazon DynamoDB

Answer: B

Explanation:

When is it appropriate to use DynamoDB instead of a relational database?

From our own experience designing and operating a highly available, highly scalable ecommerce platform, we have come to realize that relational databases should only be used when an application really needs the complex query, table join and transaction capabilities of a full-blown relational database. In all other cases, when such relational features are not needed, a NoSQL database service like DynamoDB offers a simpler, more available, more scalable and ultimately a lower cost solution.

QUESTION 281

A t2.medium EC2 instance type must be launched with what type of Amazon Machine Image (AMI)?

- A. An Instance store Hardware Virtual Machine AMI
- B. An Instance store Paravirtual AMI
- C. An Amazon EBS-backed Hardware Virtual Machine AMI
- D. An Amazon EBS-backed Paravirtual AMI

Answer: C

Explanation:

You must launch a T2 instance using an HVM AMI. For more information, see Linux AMI Virtualization Types. You must launch your T2 instances using an EBS volume as the root device. For more information, see Amazon EC2 Root Device Volume.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html>

QUESTION 282

You manually launch a NAT AMI in a public subnet. The network is properly configured. Security groups and network access control lists are properly configured. Instances in a private subnet can access the NAT. The NAT can access the Internet. However, private instances cannot access the Internet. What additional step is required to allow access from the private instances?

- A. Enable Source/Destination Check on the private Instances.
- B. Enable Source/Destination Check on the NAT instance.
- C. Disable Source/Destination Check on the private instances.
- D. Disable Source/Destination Check on the NAT instance.

Answer: D

Explanation:

Disabling Source/Destination Checks.

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. You can disable the SrcDestCheck attribute for a NAT instance that's either running or stopped using the console or the command line.

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html

QUESTION 283

Which of the following approaches provides the lowest cost for Amazon Elastic Block Store snapshots while giving you the ability to fully restore data?

- A. Maintain two snapshots: the original snapshot and the latest incremental snapshot.
- B. Maintain a volume snapshot; subsequent snapshots will overwrite one another.
- C. Maintain a single snapshot the latest snapshot is both Incremental and complete.
- D. Maintain the most current snapshot, archive the original and incremental to Amazon Glacier.

Answer: C

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-deleting-snapshot.html>

QUESTION 284

An existing application stores sensitive information on a non-boot Amazon EBS data volume attached to an Amazon Elastic Compute Cloud instance. Which of the following approaches would protect the sensitive data on an Amazon EBS volume?

- A. Upload your customer keys to AWS CloudHSM. Associate the Amazon EBS volume with AWS CloudHSM. Re-mount the Amazon EBS volume.
- B. Create and mount a new, encrypted Amazon EBS volume. Move the data to the new volume. Delete the old Amazon EBS volume.
- C. Unmount the EBS volume. Toggle the encryption attribute to True. Re-mount the Amazon EBS volume.
- D. Snapshot the current Amazon EBS volume. Restore the snapshot to a new, encrypted Amazon EBS volume. Mount the Amazon EBS volume.

Answer: B

Explanation:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html>

To migrate data between encrypted and unencrypted volumes:

1. Create your destination volume (encrypted or unencrypted, depending on your need) by following the procedures in Creating an Amazon EBS Volume.
2. Attach the destination volume to the instance that hosts the data to migrate. For more information, see Attaching an Amazon EBS Volume to an Instance.

3. Make the destination volume available by following the procedures in Making an Amazon EBS Volume Available for Using. For Linux instances, you can create a mount point at /mnt/destination and mount the destination volume there.
4. Copy the data from your source directory to the destination volume. It may be most convenient to use a bulk-copy utility for this.

QUESTION 285

A US-based company is expanding their web presence into Europe. The company wants to extend their AWS infrastructure from Northern Virginia (us-east-1) into the Dublin (eu-west-1) region. Which of the following options would enable an equivalent experience for users on both continents?

- A. Use a public-facing load balancer per region to load-balance web traffic, and enable HTTP health checks.
- B. Use a public-facing load balancer per region to load-balance web traffic, and enable sticky sessions.
- C. Use Amazon Route 53, and apply a geolocation routing policy to distribute traffic across both regions.
- D. Use Amazon Route 53, and apply a weighted routing policy to distribute traffic across both regions.

Answer: C

Explanation:

Geolocation routing lets you choose the resources that serve your traffic based on the geographic location of your users, meaning the location from which DNS queries originate. For example, you might want all queries from Africa to be routed to a web server with an IP address of 192.0.2.111.

Another possible use is for balancing load across endpoints in a predictable, easy-to-manage way, so that each user location is consistently routed to the same endpoint.

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-weighted>

QUESTION 286

Which of the following are use cases for Amazon DynamoDB? Choose 3 answers.

- A. Storing BLOB data.
- B. Managing web sessions.
- C. Storing JSON documents.
- D. Storing metadata for Amazon S3 objects.
- E. Running relational joins and complex updates.
- F. Storing large amounts of infrequently accessed data.

Answer: BCD

Explanation:

Ideal Usage Patterns.

Amazon DynamoDB is ideal for existing or new applications that need a flexible NoSQL database with low read and write latencies, and the ability to scale storage and throughput up or down as needed without code changes or downtime. Use cases require a highly available and scalable database because downtime or performance degradation has an immediate negative impact on an organization's business. for e.g. mobile apps, gaming, digital ad serving, live voting and audience interaction for live events, sensor networks, log ingestion, access control for web-based content, metadata storage for Amazon S3 objects, e-commerce shopping carts, and web session management.

QUESTION 287

A customer implemented AWS Storage Gateway with a gateway-cached volume at their main office. An event takes the link between the main and branch office offline. Which methods will enable the branch office to access their data? Choose 3 answers.

- A. Use a HTTPS GET to the Amazon S3 bucket where the files are located.
- B. Restore by implementing a lifecycle policy on the Amazon S3 bucket.
- C. Make an Amazon Glacier Restore API call to load the files into another Amazon S3 bucket within four to six hours.
- D. Launch a new AWS Storage Gateway instance AMI in Amazon EC2, and restore from a gateway snapshot.
- E. Create an Amazon EBS volume from a gateway snapshot, and mount it to an Amazon EC2 instance.
- F. Launch an AWS Storage Gateway virtual iSCSI device at the branch office, and restore from a gateway snapshot.

Answer: DEF

Explanation:

A is certainly not right, because files persisted by Storage Gateway to S3 are not visible, let alone be accessible. <https://forums.aws.amazon.com/thread.jspa?threadID=109748>

B is invalid option because you cannot apply Lifecycle Policies because AWS Storage Gateway does not give you that option. Cached Volumes are never stored to Glacier and hence "C" is not a valid.

QUESTION 288

A company has configured and peered two VPCs: VPC-1 and VPC-2. VPC-1 contains only private subnets, and VPC-2 contains only public subnets. The company uses a single AWS Direct Connect connection and private virtual interface to connect their on-premises network with VPC-1. Which two methods increase the fault tolerance of the connection to VPC-1? Choose 2 answers

- A. Establish a hardware VPN over the internet between VPC-2 and the on-premises network.
- B. Establish a hardware VPN over the internet between VPC-1 and the on-premises network.
- C. Establish a new AWS Direct Connect connection and private virtual interface in the same region as VPC-2.
- D. Establish a new AWS Direct Connect connection and private virtual interface in a different AWS region than VPC-1.
- E. Establish a new AWS Direct Connect connection and private virtual interface in the same AWS region as VPC-1.

Answer: BE

QUESTION 289

What is the minimum time Interval for the data that Amazon CloudWatch receives and aggregates?

- A. One second
- B. Five seconds
- C. One minute
- D. Three minutes
- E. Five minutes

Answer: C

Explanation:

Many metrics are received and aggregated at 1-minute intervals. Some are at 3-minute or 5-minute intervals.

QUESTION 290

Which of the following statements are true about Amazon Route 53 resource records? Choose 2 answers.

- A. An Alias record can map one DNS name to another Amazon Route 53 DNS name.
- B. A CNAME record can be created for your zone apex.
- C. An Amazon Route 53 CNAME record can point to any DNS record hosted anywhere.
- D. TTL can be set for an Alias record in Amazon Route 53.
- E. An Amazon Route 53 Alias record can point to any DNS record hosted anywhere.

Answer: AC

Explanation:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

QUESTION 291

A 3-tier e-commerce web application is currently deployed on-premises and will be migrated to AWS for greater scalability and elasticity. The web server currently shares read-only data using a network distributed file system. The app server tier uses a clustering mechanism for discovery and shared session state that depends on IP multicast. The database tier uses shared-storage clustering to provide database fail-over capability, and uses several read slaves for scaling. Data on all servers and the distributed file system directory is backed up weekly to off-site tapes. Which AWS storage and database architecture meets the requirements of the application?

- A. Web servers store read-only data in S3, and copy from S3 to root volume at boot time. App servers share state using a combination of DynamoDB and IP unicast. Database uses RDS with multi-AZ deployment and one or more Read Replicas. Backup web and app servers backed up weekly via Mils database backed up via DB snapshots.
- B. Web servers store read-only data in S3, and copy from S3 to root volume at boot time. App servers share state using a combination of DynamoDB and IP unicast. Database uses RDS with multi-AZ deployment and one or more read replicas. Backup web servers, app servers, and database backed up weekly to Glacier using snapshots.
- C. Web servers store read-only data in S3 and copy from S3 to root volume at boot time. App servers share state using a combination of DynamoDB and IP unicast. Database uses RDS with multi-AZ deployment. Backup web and app servers backed up weekly via AM is. Database backed up via DB snapshots.
- D. Web servers store read-only data in an EC2 NFS server, mount to each web server at boot time. App servers share state using a combination of DynamoDB and IP multicast. Database uses RDS with multi-AZ deployment and one or more Read Replicas. Backup web and app servers backed up weekly via Mils database backed up via DB snapshots.

Answer: A

Explanation:

<https://d0.awsstatic.com/whitepapers/Storage/AWS%20Storage%20Services%20Whitepaper-v9.pdf>

Amazon Glacier doesn't suit all storage situations. Listed following are a few storage needs for which you should consider other AWS storage options instead of Amazon Glacier.

Data that must be updated very frequently might be better served by a storage solution with lower read/write latencies, such as Amazon EBS, Amazon RDS, Amazon DynamoDB, or relational databases running on EC2.

QUESTION 292

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

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

Answer: A

Explanation:

You need to use enterprise backup software to provide file level restore. See:

https://d0.awsstatic.com/whitepapers/Backup_and_Recovery_Approaches_Using_AWS.pdf

Page 18:

If your existing backup software does not natively support the AWS cloud, you can use AWS storage gateway products. AWS Storage Gateway is a virtual appliance that provides seamless and secure integration between your data center and the AWS storage infrastructure.

QUESTION 293

Your company has HQ in Tokyo and branch offices all over the world and is using a logistics software with a multi-regional deployment on AWS in Japan, Europe and USA. The logistic software has a 3-tier architecture and currently uses MySQL 5.6 for data persistence. Each region has deployed its own database. In the HQ region you run an hourly batch process reading data from every region to compute cross-regional reports that are sent by email to all offices this batch process must be completed as fast as possible to quickly optimize logistics how do you build the database architecture in order to meet the requirements?

- A. For each regional deployment, use RDS MySQL with a master in the region and a read replica in the HQ region.
- B. For each regional deployment, use MySQL on EC2 with a master in the region and send hourly EBS snapshots to the HQ region.
- C. For each regional deployment, use RDS MySQL with a master in the region and send hourly RDS snapshots to the HQ region.
- D. For each regional deployment, use MySQL on EC2 with a master in the region and use S3 to copy data files hourly to the HQ region.
- E. Use Direct Connect to connect all regional MySQL deployments to the HQ region and reduce network latency for the batch process.

Answer: A

QUESTION 294

A customer has a 10 GB AWS Direct Connect connection to an AWS region where they have a web application hosted on Amazon Elastic Computer Cloud (EC2). The application has dependencies on an on-premises mainframe database that uses a BASE (Basic Available. Sort stale Eventual consistency) rather than an ACID (Atomicity. Consistency isolation. Durability) consistency model. The application is exhibiting undesirable behavior because the database is not able to handle the volume of writes. How can you reduce the load on your on-premises database resources in the most cost-effective way?

- A. Use an Amazon Elastic Map Reduce (EMR) S3DistCp as a synchronization mechanism between the on-premises database and a Hadoop cluster on AWS.

- B. Modify the application to write to an Amazon SQS queue and develop a worker process to flush the queue to the on-premises database.
- C. Modify the application to use DynamoDB to feed an EMR cluster which uses a map function to write to the on-premises database.
- D. Provision an RDS read-replica database on AWS to handle the writes and synchronize the two databases using Data Pipeline.

Answer: B

Explanation:

<https://aws.amazon.com/sqs/faqs/>

QUESTION 295

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

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

Answer: B

Explanation:

The requirements state “Users will log into the game using their existing social media account to streamline data capture.” This is what Cognito is used for, i.e. Web Identity Federation. Amazon also recommends to “build your app so that it requests temporary AWS security credentials dynamically when needed using web identity federation.”

QUESTION 296

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

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

Answer: B

Explanation:

<https://aws.amazon.com/sqs/faqs/>

There is no limit on the number of messages that can be pushed onto SQS. The retention period of the SQS is 4 days by default and it can be changed to 14 days. This will make sure that no writes are missed.

QUESTION 297

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

- A. Larger storage volumes support higher Provisioned IOPS rates: increase the provisioned volume storage of each of the 6 EBS volumes to 1TB.
- B. The EBS-Optimized throughput limits the total IOPS that can be utilized: use an EBS-Optimized instance that provides larger throughput.

- C. Small block sizes cause performance degradation, limiting the I/O throughput, configure the instance device driver and file system to use 64KB blocks to increase throughput.
- D. RAID 0 only scales linearly to about 4 devices, use RAID 0 with 4 EBS Provisioned IOPS volumes but increase each Provisioned IOPS EBS volume to 6,000 IOPS.
- E. The standard EBS instance root volume limits the total IOPS rate, change the instant root volume to also be a 500GB 4,000 Provisioned IOPS volume.

Answer: E

QUESTION 298

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 of 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 of 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 will 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: B

Explanation:

The POC solution is being scaled up by 1000, which means it will require 72TB of Storage to retain 24 months' worth of data. This rules out RDS as a possible DB solution which leaves you with RedShift. I believe DynamoDB is a more cost effective and scales better for ingest rather than using EC2 in an auto scaling group. Also, this example solution from AWS is somewhat similar for reference.

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_timeseriesprocessing_16.pdf

QUESTION 299

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets. Each collar will push 30kb of biometric data in JSON format every 2 seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal. Management has tasked you to architect the collection platform ensuring the following requirements are met. Provide the ability for real-time analytics of the inbound biometric data. Ensure processing of the biometric data is highly durable. Elastic and parallel. The results of the analytic processing should be persisted for data mining. Which architecture outlined below will meet the initial requirements for the collection platform?

- A. Utilize S3 to collect the inbound sensor data, analyze the data from S3 with a daily scheduled Data Pipeline and save the results to a Redshift Cluster.
- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data, analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EMR to collect the inbound sensor data, analyze the data from EMR with Amazon Kinesis and save the results to DynamoDB.

Answer: B

QUESTION 300

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

- A. Use RDS Multi-AZ with two tables, one for "Active calls" and one for "Terminated calls". In this way the "Active calls_"

table is always small and effective to access.

- B. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only. In this way the Global Secondary index is sparse and more effective.
- C. Use DynamoDB with a "Calls" table and a Global secondary index on a "State" attribute that can equal to "active" or "terminated" in this way the Global Secondary index can be used for all items in the table.
- D. Use RDS Multi-AZ with a "CALLS" table and an Indexed "STATE" field that can be equal to "ACTIVE" or "TERMINATED" in this way the SQL query is optimized by the use of the Index.

Answer: B

Explanation:

<https://aws.amazon.com/dynamodb/faqs/>

Q: Can a global secondary index key be defined on non-unique attributes?

Yes. Unlike the primary key on a table, a GSI index does not require the indexed attributes to be unique.

Q: Are GSI key attributes required in all items of a DynamoDB table?

No. GSIs are sparse indexes. Unlike the requirement of having a primary key, an item in a DynamoDB table does not have to contain any of the GSI keys. If a GSI key has both hash and range elements, and a table item omits either of them, then that item will not be indexed by the corresponding GSI. In such cases, a GSI can be very useful in efficiently locating items that have an uncommon attribute.

➤ **Vendor: Amazon**

➤ **Exam Code: AWS Certified Solutions Architect - Associate**

➤ **Exam Name: AWS Certified Solutions Architect - Associate**

➤ **Question 301 – End**

QUESTION 301

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

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

Answer: C

Explanation:

In question we have keywords 'scalable' and company wants to 'move systems' to AWS, which is best suited for Auto-scaling group.

QUESTION 302

A group can contain many users. Can a user belong to multiple groups?

- A. Yes always
- B. No
- C. Yes but only if they are using two factor authentication
- D. Yes but only in VPC

Answer: A

Explanation:

A group can contain many users, and a user can belong to multiple groups.

http://docs.aws.amazon.com/IAM/latest/UserGuide/id_groups.html

QUESTION 303

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: AB

Explanation:

As per the document defined, new IPs should be reserved not the same ones. Elastic IP Addresses are static IP addresses designed for dynamic cloud computing. Unlike traditional static IP addresses, however, Elastic IP addresses enable you to mask instance or Availability Zone failures by programmatically remapping your public IP addresses to instances in your account in a particular region. For DR, you can also pre-allocate some IP addresses for the most critical systems so that their IP addresses are already known before disaster strikes. This can simplify the execution of the DR plan.

QUESTION 304

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 servers and 6 application tier servers for normal operation, but can run on a minimum of 65% server capacity and a single MySQL 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 (Relational Database 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. d 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 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

Explanation:

<http://highscalability.com/blog/2016/1/11/a-beginners-guide-to-scaling-to-11-million-users-on-amazons.html>

<https://www.airpair.com/aws/posts/building-a-scalable-web-app-on-amazon-web-services-p1>

QUESTION 305

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

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

Answer: B

Explanation:

Answer A is mentioning 2 clusters not 1 cluster (with 2 nodes). If a node fails due to a hardware fault in an underlying physical server, ElastiCache will provision a new node on a different server.

QUESTION 306

You are responsible for a legacy web application whose server environment is approaching end of life You would like to migrate this application to AWS as quickly as possible, since the application environment currently has the following

limitations:

- The VM's single 10GB VMDK is almost full.
 - The virtual network interface still uses the 10Mbps 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 the installation media. This is a mission critical application with an RTO (Recovery Time Objective) of 8 hours. RPO (Recovery Point Objective) of 1 hour. How could you best migrate this application to AWS while meeting your business continuity requirements?

- A. Use the EC2 VM Import Connector for vCenter to import the VM into EC2.
- B. Use Import/Export to import the VM 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 the ec2-bundle-instance API to import an image of the VM into EC2.

Answer: A

Explanation:

<https://aws.amazon.com/developertools/2759763385083070>

QUESTION 307

An International company has deployed a multi-tier web application that relies on DynamoDB in a single region. For regulatory reasons they need disaster recovery capability in a separate region with a Recovery Time Objective of 2 hours and a Recovery Point Objective of 24 hours. They should synchronize their data on a regular basis and be able to provision the web application rapidly using CloudFormation. The objective is to minimize changes to the existing web application, control the throughput of DynamoDB used for the synchronization of data and synchronize only the modified elements. Which design would you choose to meet these requirements?

- A. Use AWS Data Pipeline to schedule a DynamoDB cross-region copy once a day. Create a 'LastUpdated' attribute in your DynamoDB table that would represent the timestamp of the last update and use it as a filter.
- B. Use EMR and write a custom script to retrieve data from DynamoDB in the current region using a SCAN operation and push it to DynamoDB in the second region.
- C. Use AWS Data Pipeline to schedule an export of the DynamoDB table to S3 in the current region once a day then schedule another task immediately after it that will import data from S3 to DynamoDB in the other region.
- D. Send each update into an SQS queue in the second region; use an auto-scaling group behind the SQS queue to replay the write in the second region.

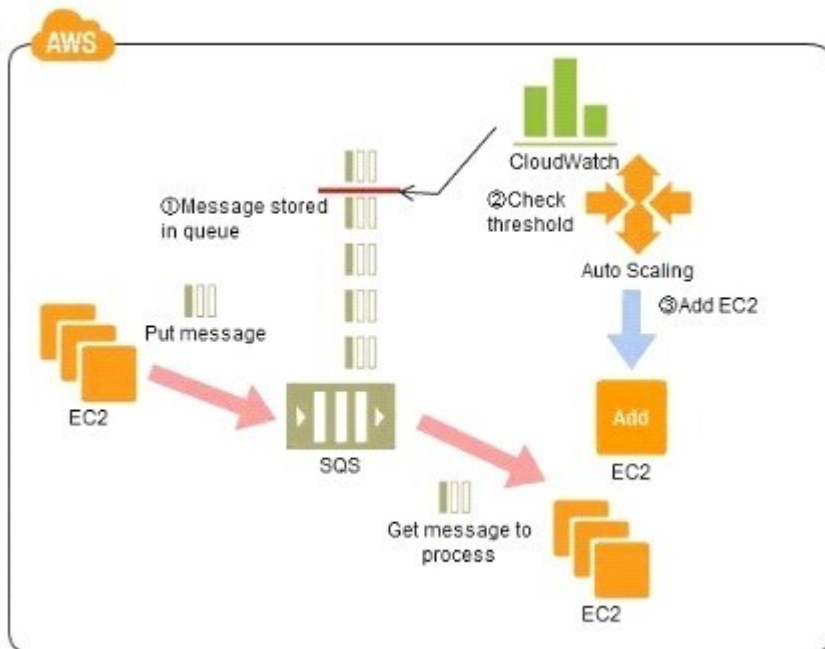
Answer: C

Explanation:

Export of DynamoDB is incremental and it will amend the backup with latest changes.

QUESTION 308

Refer to the architecture diagram above of a batch processing solution using Simple Queue Service (SQS) to set up a message queue between EC2 instances which are used as batch processors. CloudWatch monitors the number of Job requests (queued messages) and an Auto Scaling group adds or deletes batch servers automatically based on parameters set in CloudWatch alarms. You can use this architecture to implement which of the following features in a cost-effective and efficient manner?



- A. Reduce the overall time for executing jobs through parallel processing by allowing a busy EC2 instance that receives a message to pass it to the next instance in a daisy-chain setup.
- B. Implement fault tolerance against EC2 instance failure since messages would remain in SQS and work can continue with recovery of EC2 instances implement fault tolerance against SQS failure by backing up messages to S3.
- C. Implement message passing between EC2 instances within a batch by exchanging messages through SOS.
- D. Coordinate number of EC2 instances with number of job requests automatically thus Improving cost effectiveness.
- E. Handle high priority jobs before lower priority jobs by assigning a priority metadata field to SQS messages.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-using-sqs-queue.html>

QUESTION 309

Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection. How would you do this while minimizing costs?

- A. Create an EBS backed private AMI which includes a fresh install of your application. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an S3 bucket using multi-part upload.
- B. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.
- C. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zones asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- D. Create an EBS backed private AMI that includes a fresh install of your application. Develop a Cloud Formation template which includes your AMI and the required EC2, Auto-Scaling and ELB resources to support deploying the application across Multiple-Availability Zones. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.

Answer: D

QUESTION 310

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 minutes.

Answer: A

QUESTION 311

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 then 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 servers and re-use the RDS 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 activity 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 activity 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 them. Use SES to send emails to customers.

Answer: C

Explanation:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_ecommerce_checkout_13.pdf

QUESTION 312

You have deployed a web application targeting a global audience across multiple AWS Regions under the domain name.example.com. You decide to use Route53 Latency-Based Routing to serve web requests to users from the region closest to the user. To provide business continuity in the event of server downtime you configure weighted record sets associated with two web servers in separate Availability Zones per region. During a DR test you notice that when you disable all web servers 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 for one or more of the weighted resource record sets associated with the disabled web servers.
- C. The value of the weight associated with the latency alias resource record set in the region with the disabled servers is higher than the weight for the other region.
- D. One of the two working web servers 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

Explanation:

<http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-complex-configs.html>

For both latency alias resource record sets, you set the value of "Evaluate Target Health" to Yes. You use the Evaluate Target Health setting for each latency alias resource record set to make Amazon Route 53 evaluate the health of the alias targets—the weighted resource record sets—and respond accordingly.

QUESTION 313

Your company hosts a social media site supporting users in multiple countries. You have been asked to provide a highly available design for the application that leverages multiple regions for the most recently accessed content and latency sensitive portions of the website. The most latency sensitive component of the application involves reading user preferences to support website 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. CloudFront 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. CloudFront with dynamic content, and an ELB in each region Retrieve user preferences from an ElasticCache cluster in each region and leverage Simple Workflow (SWF) to manage the propagation of user preferences from a centralized OB to each ElasticCache cluster.

Answer: A

Explanation:

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_mediasharing_09.pdf

http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_adserving_06.pdf

QUESTION 314

Your system recently experienced down time during the troubleshooting process. You found that a new administrator mistakenly terminated several production EC2 instances. Which of the following strategies will help prevent a similar situation in the future? The administrator still must be able to:

- launch, start stop, and terminate development resources.
- launch and start production instances.

- A. Create an IAM user, which is not allowed to terminate instances by leveraging production EC2 termination protection.
- B. Leverage resource based tagging along with an IAM user, which can prevent specific users from terminating production EC2 resources.
- C. Leverage EC2 termination protection and multi-factor authentication, which together require users to authenticate before terminating EC2 instances.
- D. Create an IAM user and apply an IAM role which prevents users from terminating production EC2 instances.

Answer: B

Explanation:

<http://blogs.aws.amazon.com/security/post/Tx29HCT3ABL7LP3/Resource-level-Permissions-for-EC2-Controlling-Management-Access-on-Specific-Ins>

QUESTION 315

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

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

Answer: BE

Explanation:

<https://myawsscribble.wordpress.com/2015/09/25/setting-up-and-configuring-aws-directconnect/>

QUESTION 316

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 verify network traffic is leveraging DirectConnect.

- B. Configure your DirectConnect router with a higher BGP priority than 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 priority. And verify network traffic is leveraging the DirectConnect connection.

Answer: C

Explanation:

Direct Connect takes priority over Dynamically configured VPN connections.

QUESTION 317

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

- A. Route payment requests through two NAT instances setup for High Availability and whitelist the Elastic IP addresses attached to the NAT instances.
- B. Whitelist the VPC Internet Gateway Public IP and route payment requests through the Internet Gateway.
- C. Whitelist the ELB IP addresses and route payment requests from the Application servers through the ELB.
- D. Automatically assign public IP addresses to the application instances in the Auto Scaling group and run a script on boot that adds each instance's public IP address to the payment validation whitelist API.

Answer: A

Explanation:

B is incorrect as you do not have insight into the public IP associated with a VPC Internet Gateway.

C is incorrect as ELB receives a public DNS name.

D would exceed the maximum of 4 whitelisting IP addresses.

QUESTION 318

You are designing the network infrastructure for an application server in Amazon VPC. Users will access all the application instances from the Internet as well as from an on-premises network. The on-premises network is connected to your VPC over an AWS Direct Connect link. How would you design routing to meet the above requirements?

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

Answer: B

QUESTION 319

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 link. Configure a static route via your AWS Direct connect link that points to Amazon S3. Configure specific routes to your network in your VPC.
- C. 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.

- D. Create a private interface on your AWS Direct connect link. Redistribute BGP routes into your existing routing infrastructure and advertise a default route to AWS.

Answer: C

Explanation:

<https://aws.amazon.com/directconnect/faqs/>

QUESTION 320

You have deployed a three-tier web application in a VPC with a CIDR block of 10.0.0.0/28. You initially deploy two web servers, two application servers, two database servers and one NAT instance for a total of seven EC2 instances. The web, application and database servers are deployed across two availability zones (AZs). You also deploy an ELB in front of the two web servers, and use Route53 for DNS. Web traffic 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 cause? Choose 2 answers.

- A. 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.
- B. 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.
- C. 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.
- D. The ELB has scaled-up. Adding more instances to handle the traffic reducing the number of available private IP addresses for new instance launches.
- E. AWS reserves the first 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: DE

Explanation:

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

QUESTION 321

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-2f8t>C447

IGW: ig-2d8bc445

NACL: acl-2080c448

Subnets and Route Tables:

Web server's subnet: subnet-258Dc44d

Application server's subnet: subnet-248bc44c

Database server's subnet: subnet-9189c6f9

Route Tables:

rtb-218DC449

rtb-238bc44b

Associations:

subnet-258bc44d: rtb-218bc449

Subnet-248DC44C: rtb-238tX44b

subnet-9189c6f9: rtb-238Dc 44b

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 servers to retrieve updates from the Internet?

- A. Create a bastion and NAT Instance in subnet-248bc44c and add a route from rtb-238bc44b to subnet-258bc44d.
- 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-258bc44d. Add a route from rtb-238bc44b to igw-2d8bc445. And a new NACL that allows access between subnet-258bc44d and subnet-248bc44c.
- D. Create a bastion and NAT instance in subnet-258bc44d and add a route from rtb-238bc44b to the NAT instance.

Answer:

[Visit PassLeader and Download Full Version AWS-Associate Exam Dumps](#)

- 1) **A company is storing an access key (access key ID and secret access key) in a text file on a custom AMI. The company uses the access key to access DynamoDB tables from instances created from the AMI. The security team has mandated a more secure solution.**

Which solution will meet the security team’s mandate?

- A. Put the access key in an S3 bucket, and retrieve the access key on boot from the instance.
- B. Pass the access key to the instances through instance user data.
- C. Obtain the access key from a key server launched in a private subnet.
- D. Create an IAM role with permissions to access the table, and launch all instances with the new role.

- 2) **A company is developing a highly available web application using stateless web servers. Which services are suitable for storing session state data? (Select TWO.)**

- A. CloudWatch
- B. DynamoDB
- C. Elastic Load Balancing
- D. ElastiCache
- E. Storage Gateway

- 3) **Company salespeople upload their sales figures daily. A Solutions Architect needs a durable storage solution for these documents that also protects against users accidentally deleting important documents.**

Which action will protect against unintended user actions?

- A. Store data in an EBS volume and create snapshots once a week.
- B. Store data in an S3 bucket and enable versioning.
- C. Store data in two S3 buckets in different AWS regions.
- D. Store data on EC2 instance storage.

- 4) **An application requires a highly available relational database with an initial storage capacity of 8 TB. The database will grow by 8 GB every day. To support expected traffic, at least eight read replicas will be required to handle database reads.**

Which option will meet these requirements?

- A. DynamoDB
- B. Amazon S3
- C. Amazon Aurora
- D. Amazon Redshift

- 5) **A Solutions Architect is designing a critical business application with a relational database that runs on an EC2 instance. It requires a single EBS volume that can support up to 16,000 IOPS.**

Which Amazon EBS volume type can meet the performance requirements of this application?

- A. EBS Provisioned IOPS SSD
- B. EBS Throughput Optimized HDD
- C. EBS General Purpose SSD
- D. EBS Cold HDD

- 6) **A web application allows customers to upload orders to an S3 bucket. The resulting Amazon S3 events trigger a Lambda function that inserts a message to an SQS queue. A single EC2 instance reads messages from the queue, processes them, and stores them in an DynamoDB table partitioned by unique order ID. Next month traffic is expected to increase by a factor of 10 and a Solutions Architect is reviewing the architecture for possible scaling problems.**

Which component is MOST likely to need re-architecting to be able to scale to accommodate the new traffic?

- A. Lambda function
- B. SQS queue
- C. EC2 instance
- D. DynamoDB table

- 7) **An application saves the logs to an S3 bucket. A user wants to keep the logs for one month for troubleshooting purposes, and then purge the logs.**

What feature will enable this?

- A. Adding a bucket policy on the S3 bucket.
- B. Configuring lifecycle configuration rules on the S3 bucket.
- C. Creating an IAM policy for the S3 bucket.
- D. Enabling CORS on the S3 bucket.

- 8) **An application running on EC2 instances processes sensitive information stored on Amazon S3. The information is accessed over the Internet. The security team is concerned that the Internet connectivity to Amazon S3 is a security risk.**

Which solution will resolve the security concern?

- A. Access the data through an Internet Gateway.
- B. Access the data through a VPN connection.
- C. Access the data through a NAT Gateway.
- D. Access the data through a VPC endpoint for Amazon S3.

- 9) **An organization is building an Amazon Redshift cluster in their shared services VPC. The cluster will host sensitive data.**

How can the organization control which networks can access the cluster?

- A. Run the cluster in a different VPC and connect through VPC peering.
- B. Create a database user inside the Amazon Redshift cluster only for users on the network.
- C. Define a cluster security group for the cluster that allows access from the allowed networks.
- D. Only allow access to networks that connect with the shared services network via VPN.

- 10) **A Solutions Architect is designing an online shopping application running in a VPC on EC2 instances behind an ELB Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application tier must read and write data to a customer managed database cluster. There should be no access to the database from the Internet, but the cluster must be able to obtain software patches from the Internet.**

Which VPC design meets these requirements?

- A. Public subnets for both the application tier and the database cluster
- B. Public subnets for the application tier, and private subnets for the database cluster
- C. Public subnets for the application tier and NAT Gateway, and private subnets for the database cluster
- D. Public subnets for the application tier, and private subnets for the database cluster and NAT Gateway

Answers

- 1) D – IAM roles for EC2 instances allow applications running on the instance to access AWS resources without having to create and store any access keys. Any solution involving the creation of an access key then introduces the complexity of managing that secret.
- 2) B, D – Both DynamoDB and ElastiCache provide high performance storage of key-value pairs. CloudWatch and ELB are not storage services. Storage Gateway is a storage service, but it is a hybrid storage service that enables on-premises applications to use cloud storage.
- 3) B – If a versioned object is deleted, then it can still be recovered by retrieving the final version. Response A would lose any changes committed since the previous snapshot. Storing the data in 2 S3 buckets would provide slightly more protection, but a user could still delete the object from both buckets. EC2 instance storage is ephemeral and should never be used for data requiring durability.
- 4) C – Amazon Aurora is a relational database that will automatically scale to accommodate data growth. Amazon Redshift does not support read replicas and will not automatically scale. DynamoDB is a NoSQL service, not a relational database. Amazon S3 is object storage, not a relational database.
- 5) A – EBS Provisioned IOPS SSD provides sustained performance for mission-critical low-latency workloads. EBS General Purpose SSD can provide bursts of performance up to 3,000 IOPS and have a maximum baseline performance of 10,000 IOPS for volume sizes greater than 3.3 TB. The 2 HDD options are lower cost, high throughput volumes.
- 6) C – A single EC2 instance will not scale and is a single point of failure in the architecture. A much better solution would be to have EC2 instances in an Auto Scaling group across 2 availability zones read messages from the queue. The other responses are all managed services that can be configured to scale or will scale automatically.
- 7) B – Lifecycle configuration allows lifecycle management of objects in a bucket. The configuration is a set of one or more rules, where each rule defines an action for Amazon S3 to apply to a group of objects. Bucket policies and IAM define access to objects in an S3 bucket. CORS enables clients in one domain to interact with resources in a different domain.
- 8) D – VPC endpoints for Amazon S3 provide secure connections to S3 buckets that do not require a gateway or NAT instances. NAT Gateways and Internet Gateways still route traffic over the Internet to the public endpoint for Amazon S3. There is no way to connect to Amazon S3 via VPN.
- 9) C – A security group can grant access to traffic from the allowed networks via the CIDR range for each network. VPC peering and VPN are connectivity services and cannot control traffic for security. Amazon Redshift user accounts address authentication and authorization at the user level and have no control over network traffic.
- 10) C – The online application must be in public subnets to allow access from clients' browsers. The database cluster must be in private subnets to meet the requirement that there be no access from the Internet. A

NAT Gateway is required to give the database cluster the ability to download patches from the Internet.
NAT Gateways must be deployed in public subnets.

AWS-SOLUTION-ARCHITECT- ASSOCIATE^{Q&As}

AWS Certified Solutions Architect - Associate

QUESTION 1

A business team requires a structured storage solution to store all of a company's historical sales data. Currently there are 4 TB of data, which will grow to hundreds of terabytes within a few years. The team must be able to regularly run queries against the data using current business intelligence tools. Fast performance is required despite the dataset growth.

Which solution should the company use?

- A. Amazon Redshift
- B. Amazon Aurora
- C. Amazon DynamoDB
- D. Amazon S3

Correct Answer: A

Reference: <https://aws.amazon.com/blogs/aws/category/amazon-redshift/>

QUESTION 2

A company plans to use an Amazon VPC to deploy a web application consisting of an elastic load balancer, a fleet of web and application servers, and an Amazon RDS MySQL database that should not be accessible from the Internet. The proposed design must be highly available and distributed over two Availability Zones.

What would be the MOST appropriate VPC design for this specific use case?

- A. Two public subnets for the elastic load balancer, two public subnets for the web servers, and two public subnets for Amazon RDS.
- B. One public subnet for the elastic load balancer, two private subnets for the web servers, and two private subnets for Amazon RDS.
- C. One public subnet for the elastic load balancer, one public subnet for the web servers, and one private subnet for the database.
- D. Two public subnets for the elastic load balancer, two private subnets for the web servers, and two private subnets for RDS.

Correct Answer: B

QUESTION 3

A company will run different data analytics jobs on large petabyte-scale datasets, using standard SQL and existing business intelligence tools. The data is mostly structured, but part of the data is unstructured and resides in Amazon S3. What technology should be used to support this use case?

- A. An Amazon Aurora database cluster with 15 replicas distributed across Availability Zones.
-

- B. Amazon Redshift with Amazon Redshift Spectrum.
- C. Amazon DynamoDB with Amazon DynamoDB Accelerator (DAX).
- D. Amazon ElastiCache for Redis with cluster mode enabled.

Correct Answer: B

For data warehousing, Amazon Redshift provides the ability to run complex, analytic queries against petabytes of structured data, and includes Redshift Spectrum that runs SQL queries directly against Exabytes of structured or unstructured data in S3 without the need for unnecessary data movement.

Reference: <https://aws.amazon.com/big-data/datalakes-and-analytics/>

QUESTION 4

A company plans to use Amazon GuardDuty to detect unexpected and potentially malicious activity. The company wants to use Amazon CloudWatch to ensure that when findings occur, remediation takes place automatically. Which CloudWatch feature should be used to trigger an AWS Lambda function to perform the remediation?

- A. Events
- B. Dashboards
- C. Metrics D. Alarms

Correct Answer: A

Reference: <https://aws.amazon.com/blogs/security/how-to-use-amazon-guardduty-and-aws-web-application-firewall-to-automatically-block-suspicious-hosts/>

QUESTION 5

A Solutions Architect designed a system based on Amazon Kinesis Data Streams. After the workflow was put into production, the company noticed it performed slowly and identified Kinesis Data Streams as the problem. One of the streams has a total of 10 Mb/s throughput.

What should the Solutions Architect recommend to improve performance?

- A. Use AWS Lambda to preprocess the data and transform the records into a simpler format, such as CSV.
- B. Run the MergeShard command to reduce the number of shards that the consumer can more easily process.
- C. Change the workflow to use Amazon Kinesis Data Firehose to gain a higher throughput.
- D. Run the UpdateShardCount command to increase the number of shards in the stream

Correct Answer: D

QUESTION 6

A Solutions Architect is developing a solution for sharing files in an organization. The solution must allow multiple users to access the storage service at once from different virtual machines and scale automatically. It must also support file-level locking.

Which storage service meets the requirements of this use case?

- A. Amazon S3
- B. Amazon EFS
- C. Amazon EBS
- D. Cached Volumes

Correct Answer: B

QUESTION 7

A Solutions Architect must select the storage type for a big data application that requires very high sequential I/O. The data must persist if the instance is stopped. Which of the following storage types will provide the best fit at the LOWEST cost for the application?

- A. An Amazon EC2 instance store local SSD volume.
- B. An Amazon EBS provisioned IOPS SSD volume.
- C. An Amazon EBS throughput optimized HDD volume.
- D. An Amazon EBS general purpose SSD volume.

Correct Answer: C

QUESTION 8

Your company is in the process of developing a next generation pet collar that collects biometric information to assist families with promoting healthy lifestyles for their pets. Each collar will push 30kb of biometric data in JSON format every 2

seconds to a collection platform that will process and analyze the data providing health trending information back to the pet owners and veterinarians via a web portal. Management has tasked you to architect the collection platform ensuring

the following requirements are met.

Provide the ability for real-time analytics of the inbound biometric data.

Ensure processing of the biometric data is highly durable, elastic, and parallel.

The results of the analytic processing should be persisted for data mining.

Which architecture outlined below will meet the initial requirements for the collection platform?

- A. Utilize S3 to collect the inbound sensor data, analyze the data from S3 with a daily scheduled Data Pipeline, and save the results to a Redshift Cluster.

- B. Utilize Amazon Kinesis to collect the inbound sensor data, analyze the data with Kinesis clients and save the results to a Redshift cluster using EMR.
- C. Utilize SQS to collect the inbound sensor data analyze the data from SQS with Amazon Kinesis and save the results to a Microsoft SQL Server RDS instance.
- D. Utilize EMR to collect the inbound sensor data, analyze the data from EUR with Amazon Kinesis and save me results to DynamoDB.

Correct Answer: B

The POC solution is being scaled up by 1000, which means it will require 72TB of Storage to retain 24 months\ worth of data. This rules out RDS as a possible DB solution which leaves you with RedShift. I believe DynamoDB is a more cost effective and scales better for ingest rather than using EC2 in an auto scaling group. Also, this example solution from AWS is somewhat similar for reference. http://media.amazonwebservices.com/architecturecenter/AWS_ac_ra_timeseriesprocessing_16.pdf

QUESTION 9

A company deployed a three-tier web application on Amazon EBS backed Amazon EC2 instances for the web and application tiers, and Amazon RDS for the database tier. The company is concerned about loss of data in the web and application tiers.

What is the MOST efficient way to prevent data loss?

- A. Create an Amazon EFS file system and run a shell script to copy the data
- B. Create an Amazon EBS snapshot using an Amazon CloudWatch Events rule
- C. Create an Amazon S3 snapshot policy to back up the Amazon EBS volumes
- D. Create a snapshot lifecycle policy that takes periodic snapshots of the Amazon EBS volumes

Correct Answer: B

Reference <https://aws.amazon.com/ebs/>

QUESTION 10

A company is storing data in an Amazon DynamoDB table and needs to take daily backups and retain them for 6 months. How should the Solutions Architect meet these requirements without impacting the production workload?

- A. Use DynamoDB replication and restore the table from the replica
- B. Use AWS Data Pipeline and create a scheduled job to back up the DynamoDB table daily
- C. Use Amazon CloudWatch Events to trigger an AWS Lambda function that makes an on-demand backup of the table
- D. Use AWS Batch to create a scheduled backup with the default template, then back up to Amazon S3 daily.

Correct Answer: D

QUESTION 11

A Solutions Architect must design an Amazon DynamoDB table to store data about customer activities. The data is used to analyze recent customer behavior, so data that is less than a week old is heavily accessed and older data is accessed infrequently. Data that is more than one month old never needs to be referenced by the application, but needs to be archived for year-end analytics.

What is the MOST cost-efficient way to meet these requirements? (Choose two.)

- A. Use DynamoDB time-to-live settings to expire items after a certain time period.
- B. Provision a higher write capacity unit to minimize the number of partitions.
- C. Create separate tables for each week's data with higher throughput for the current week.
- D. Pre-process data to consolidate multiple records to minimize write operations.
- E. Export the old table data from DynamoDB to Amazon S3 using AWS Data Pipeline, and delete the old table.

Correct Answer: BC

QUESTION 12

An organization regularly backs up their application data. The application backups are required to be stored on Amazon S3 for a certain amount of time. The backups should be accessed instantly in the event of a disaster recovery. Which of the following Amazon S3 storage classes would be the MOST cost-effective option to meet the needs of this scenario?

- A. Glacier Storage Class
- B. Standard Storage Class
- C. Standard - Infrequent Access (IA)
- D. Reduced Redundancy Class (RRS)

Correct Answer: C

Reference: <https://aws.amazon.com/s3/features/>

QUESTION 13

A company must collect temperature data from thousands of remote weather devices. The company must also store this data in a data warehouse to run aggregations and visualizations.

Which services will meet these requirements? (Choose two.)

- A. Amazon Kinesis Data Firehouse
 - B. Amazon SQS
 - C. Amazon Redshift
 - D. Amazon SNS
-

E. Amazon DynamoDB

Correct Answer: AB

QUESTION 14

A Solutions Architect is designing a new architecture that will use an Amazon EC2 Auto Scaling group. Which of the following factors determine the health check grace period? (Select TWO.)

- A. How frequently the Auto Scaling group scales up or down.
- B. How many Amazon CloudWatch alarms are configured for status checks.
- C. How much of the application code is embedded in the AMI.
- D. How long it takes for the Auto Scaling group to detect a failure.
- E. How long the bootstrap script takes to run.

Correct Answer: AD

QUESTION 15

An application that runs on an Amazon EC2 instance must make secure calls to Amazon S3 buckets. Which steps can a Solutions Architect take to ensure that the calls are made without exposing credentials?

- A. Generate an access key ID and a secret key, and assign an IAM role with least privilege.
- B. Create an IAM policy granting access to all services and assign it to the Amazon EC2 instance profile.
- C. Create an IAM role granting least privilege and assign it to the Amazon EC2 instance profile.
- D. Generate temporary access keys to grant users temporary access to the Amazon EC2 instance.

Correct Answer: C

QUESTION 16

A web application runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. Every night, the Auto Scaling group doubles in size. Traffic

analysis shows that users in a particular region are requesting the same static content stored locally on the EC2 instances.

How can a Solutions Architect reduce the need to scale and improve application performance for the users?

- A. Re-deploy the application in a new VPC that is closer to the users making the requests.
- B. Create an Amazon CloudFront distribution for the site and redirect user traffic to the distribution.
- C. Store the contents on Amazon EFS instead of the EC2 root volume.

D. Implement Amazon Redshift to create a repository of the content closer to the users.

Correct Answer: B

QUESTION 17

A media company has more than 100TB of data to be stored and retrieved infrequently. However, the company occasionally receives requests for data within an hour. The company needs a low-cost retrieval method to handle the requests. Which service meets this requirement?

- A. Amazon S3 Standard
- B. Amazon Glacier standard retrievals
- C. Amazon Glacier bulk retrievals
- D. Amazon S3 Standard Infrequent Access

Correct Answer: D

Reference: <https://aws.amazon.com/blogs/aws/aws-storage-update-s3-glacier-price-reductions/>

QUESTION 18

You are designing Internet connectivity for your VPC. The Web servers must be available on the Internet. The application must have a highly available architecture. Which alternatives should you consider? (Choose two.)

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

Correct Answer: CD

QUESTION 19

An application is running on an Amazon EC2 instance in a private subnet. The application needs to read and write data onto Amazon Kinesis Data Streams, and corporate policy requires that this traffic should not go to the internet. How can these requirements be met?

- A. Configure a NAT gateway in a public subnet and route all traffic to Amazon Kinesis through the NAT gateway.
 - B. Configure a gateway VPC endpoint for Kinesis and route all traffic to Kinesis through the gateway VPC endpoint.
-

C. Configure an interface VPC endpoint for Kinesis and route all traffic to Kinesis through the gateway VPC endpoint.

D. Configure an AWS Direct Connect private virtual interface for Kinesis and route all traffic to Kinesis through the virtual interface.

Correct Answer: C

QUESTION 20

Which Amazon service can I use to define a virtual network that closely resembles a traditional data center?

A. Amazon VPC

B. Amazon ServiceBus

C. Amazon EMR

D. Amazon RDS

Correct Answer: A

QUESTION 21

An environment has an Auto Scaling group across two Availability Zones referred to as AZ-a and AZ-b and a default termination policy. AZ-a has four Amazon EC2 instances, and AZ-b has three EC2 instances. None of the instances is protected from a scale-in.

How will Auto Scaling proceed if there is a scale-in event?

A. Auto Scaling selects an instance to terminate randomly.

B. Auto Scaling terminates the instance with the oldest launch configuration of all instances.

C. Auto Scaling selects the Availability Zone with four EC2 instances and then continues to evaluate.

D. Auto Scaling terminates the instance with the closest next billing hour of all instances.

Correct Answer: C

For Auto Scaling Groups with Multiple Instance Types and Purchase Options, the scale-in action respects the On-Demand base amount and the percentages set for On-Demand and Spot Instances. When the group scales in, the Auto Scaling group first identifies which of these two types of instances it will terminate. Then, it chooses the Availability Zone with the most instances of that type, and applies the default or customized termination policy.

Reference: <https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-instance-termination.html>

QUESTION 22

A Solutions Architect is designing a multi-tier application consisting of an Application Load Balancer, an Amazon RDS database instance, and an Auto Scaling group on Amazon EC2 instances. Each tier is in a separate subnet. There are some EC2 instances in the subnet that belong to another application. The RDS database instance should accept traffic

only from the EC2 instances in the Auto Scaling group.

What should be done to meet these requirements?

- A. Configure the inbound network ACLs on the database subnet to accept traffic from the IP addresses of the EC2 instances only.
- B. Configure the inbound rules on the security group associated with the RDS database instance. Set the source to the security group associated with instances in the Auto Scaling group.
- C. Configure the outbound rules on the security group associated with the Auto Scaling group. Set the destination to the security group associated with the RDS database instance.
- D. Configure the inbound network ACLs on the database subnet to accept traffic only from the CIDR range of the subnet used by the Auto Scaling group.

Correct Answer: C

QUESTION 23

An on-premises database is experiencing significant performance problems when running SQL queries. With 10 users, the lookups are performing as expected. As the number of users increases, the lookups take three times longer than expected to return values to an application.

Which action should a Solutions Architect take to maintain performance as the user count increases?

- A. Use Amazon SQS.
- B. Deploy Multi-AZ RDS MySQL
- C. Configure Amazon RDS with additional read replicas.
- D. Migrate from MySQL to RDS Microsoft SQL Server.

Correct Answer: B

QUESTION 24

An organization runs an online media site, hosted on-premises. An employee posted a product review that contained videos and pictures. The review went viral and the organization needs to handle the resulting spike in website traffic. What action would provide an immediate solution?

- A. Redesign the website to use Amazon API Gateway, and use AWS Lambda to deliver content.
- B. Add server instances using Amazon EC2 and use Amazon Route 53 with a failover routing policy.
- C. Serve the images and videos via an Amazon CloudFront distribution created using the news site as the origin.
- D. Use Amazon ElasticCache for Redis for caching and reducing the load requests from the origin.

Correct Answer: C

QUESTION 25

An organization uses Amazon S3 to store video content served via its website. It only has rights to deliver this content to users within its own country and needs to restrict access. How can the organization ensure that these files are only accessible from within its country?

- A. Use a custom Amazon S3 bucket policy to allow access only to users inside the organization's country
- B. Use Amazon CloudFront and Geo Restriction to allow access only to users inside the organization's country
- C. Use an Amazon S3 bucket ACL to allow access only to users inside the organization's country
- D. Use file-based ACL permissions on each video file to allow access only to users inside the organization's country

Correct Answer: A

Reference <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>

QUESTION 26

A company is creating a web application that will run on an Amazon EC2 instance. The application on the instance needs access to an Amazon DynamoDB table for storage. What should be done to meet these requirements?

- A. Create another AWS account root user with permissions to the DynamoDB table.
- B. Create an IAM role and assign the role to the EC2 instance with permissions to the DynamoDB table.
- C. Create an identity provider and assign the identity provider to the EC2 instance with permissions to the DynamoDB table.
- D. Create identity federation with permissions to the DynamoDB table.

Correct Answer: B

QUESTION 27

A company is running its application in a single region on Amazon EC2 with Amazon EBS and Amazon S3 part of the storage design. What should be done to reduce data transfer costs?

- A. Create a copy of the compute environment in another region
- B. Convert the application to run on Lambda@Edge
- C. Create an Amazon CloudFront distribution with Amazon S3 as the origin
- D. Replicate Amazon S3 data to buckets in regions closer to the requester

Correct Answer: D

QUESTION 28

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

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

Correct Answer: C

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

QUESTION 29

A Solution Architect is creating a serverless web application that must access mapping data in hundreds of data files, each containing approximately 30 KB of data. The storage required is expected to grow to hundreds of terabytes. Which storage solution is most cost-effective, yet still meets the requirements for this use case?

- A. Amazon EFS
- B. Amazon EBS Cold HDD (sc1)
- C. Amazon S3 Standard
- D. Amazon DynamoDB

Correct Answer: D

Reference <https://medium.freecodecamp.org/how-to-create-a-serverless-service-in-15-minutes-b63af8c892e5>

QUESTION 30

A Solutions Architect is designing a service that must have four Amazon EC2 instances running between 8 AM and 6 PM daily. The service requires one EC2 instance outside of those hours. What is the MOST cost-effective way to provide enough compute?

- A. Use one Amazon EC2 Reserved Instance and use an Auto Scaling group to add and remove EC2 instances based on CPU utilization.
- B. Use one Amazon EC2 On-Demand instance and use an Auto Scaling group to add and remove EC2 instances based on CPU utilization.
- C. Use one Amazon EC2 On-Demand instance and use an Auto Scaling Group scheduled action to add three EC2 Spot instances at 7:30 AM and remove three instances at 6:10 PM.
- D. Use one Amazon EC2 Reserved Instance and use an Auto Scaling Group scheduled action to add three EC2 On-Demand instances at 7:30 AM and remove three instances at 6:10 PM.

Correct Answer: B

QUESTION 31

Amazon rds provides a facility to modify the back-up retention policy for automated backups, with a value of 0 indicating for no backup retention. What is the maximum retention period allowed in days?

- A. 45
- B. 35
- C. 15
- D. 10

Correct Answer: B

QUESTION 32

An Internet-facing multi-tier web application must be highly available. An ELB Classic Load Balancer is deployed in front of the web tier. Amazon EC2 instances at the web application tier are deployed evenly across two Availability Zones. The

database is deployed using RDS Multi-AZ. A NAT instance is launched for Amazon EC2 instances and database resources to access the Internet. These instances are not assigned with public IP addresses.

Which component poses a potential single point of failure in this architecture?

- A. Amazon EC2
- B. NAT instance
- C. ELB Classic Load Balancer
- D. Amazon RDS

Correct Answer: B

QUESTION 33

A Solutions Architect is designing a public-facing web application for employees to upload images to their social media account. The application consists of multiple Amazon EC2 instances behind an elastic load balancer, an Amazon S3 bucket where uploaded images are stored, and an Amazon DynamoDB table for storing image metadata.

Which AWS service can the Architect use to automate the process of updating metadata in the DynamoDB table upon image upload?

- A. Amazon CloudWatch
- B. AWS CloudFormation

- C. AWS Lambda
- D. Amazon SQS

Correct Answer: C

QUESTION 34

A company is developing several critical long-running applications hosted on Docker.

How should a Solutions Architect design a solution to meet the scalability and orchestration requirements on AWS?

- A. Use Amazon ECS and Service Auto Scaling.
- B. Use Spot Instances for orchestration and for scaling containers on existing Amazon EC2 instances.
- C. Use AWS OpsWorks to launch containers in new Amazon EC2 instances.
- D. Use Auto Scaling groups to launch containers on existing Amazon EC2 instances.

Correct Answer: A

Reference: <https://aws.amazon.com/getting-started/tutorials/deploy-docker-containers/>

QUESTION 35

A Solutions Architect is designing a microservice to process records from Amazon Kinesis Streams. The metadata must be stored in Amazon DynamoDB. The microservice must be capable of concurrently processing 10,000 records daily as they arrive in the Kinesis stream.

The MOST scalable way to design the microservice is:

- A. As an AWS Lambda function.
- B. As a process on an Amazon EC2 instance.
- C. As a Docker container running on Amazon ECS.
- D. As a Docker container on an EC2 instance.

Correct Answer: C

Reference: <https://aws.amazon.com/blogs/big-data/optimize-spark-streaming-to-efficiently-process-amazon-kinesis-streams/>

QUESTION 36

What does a "Domain" refer to in Amazon SWF?

- A. A security group in which only tasks inside can communicate with each other

- B. A special type of worker
- C. A collection of related Workflows
- D. The DNS record for the Amazon SWF service

Correct Answer: C

Domains provide a way of scoping Amazon SWF resources within your AWS account. All the components of a workflow, such as the workflow type and activity types, must be specified to be in a domain. It is possible to have more than one workflow in a domain; however, workflows in different domains cannot interact with each other.
<http://docs.aws.amazon.com/amazonswf/latest/developerguide/swf-dev-domain.html>

QUESTION 37

An application runs on EC2 instances behind an Elastic Load Balancing Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The application provides a RESTful interface with both synchronous and asynchronous operations. The asynchronous operations require up to 5 minutes to complete. Although the application must remain available at all times, after business hours, the traffic going to the application is greatly reduced and often results in the Auto Scaling group running the minimum number of On-Demand Instances.

What should the Solutions Architect recommend to optimize the cost of the environment after business hours?

- A. Change the Availability Zones in which the instances were created to another Availability Zone in the same region with a lower cost.
- B. Replace all On-Demand Instances with Spot Instances in the Auto Scaling group.
- C. Purchase Reserved Instances for the minimum number of Auto Scaling instances.
- D. Reduce the number of minimum instances to 0. New requests to the Application Load Balancer create new instances.

Correct Answer: C

QUESTION 38

A Solutions Architect is creating a new relational database. The Compliance team will use the database, and mandates that data content must be stored across three different Availability Zones. Which of the following options should the Architect Use?

- A. Amazon Aurora
- B. Amazon RDS MySQL with Multi-AZ enabled
- C. Amazon DynamoDB
- D. Amazon ElastiCache

Correct Answer: B

QUESTION 39

Does Amazon RDS support SSL encryption for SQL Server DB Instances?

- A. Yes, for all supported SQL Server editions
- B. No
- C. Yes, but only when the instances are in a single region
- D. No, encryption using SSL is supported only in the GovCloud.

Correct Answer: A

Amazon RDS supports SSL encryption for SQL Server DB Instances. Using SSL, you can encrypt connections between your applications and your SQL Server DB Instances. This is available for all the versions of Microsoft SQL Server. <http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.DBEngine.SQLServer.html>

QUESTION 40

A company's policy requires that all data stored in Amazon S3 is encrypted. The company wants to use the option with the least overhead and does not want to manage any encryption keys. Which of the following options will meet the company's requirements?

- A. AWS CloudHSM
- B. AWS Trusted Advisor
- C. Server Side Encryption (SSE-S3)
- D. Server Side Encryption (SSE-KMS)

Correct Answer: C

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/dev/bucket-encryption.html>

AWS-SOLUTION-ARCHITECT- ASSOCIATE^{Q&As}

AWS Certified Solutions Architect - Associate

QUESTION 1

Which of the following cannot be used in Amazon EC2 to control who has access to specific Amazon EC2 instances?

- A. Security Groups
- B. IAM System
- C. SSH keys
- D. Windows passwords

Correct Answer: B

<http://blogs.aws.amazon.com/security/post/Tx29HCT3ABL7LP3/Resource-level-Permissions-for-EC2-Controlling-Management-Access-on-Specific-Ins>

QUESTION 2

A customer is deploying a production portal application on AWS. The database tier has structured data. The company requires a solution that is easily manageable and highly available. How can these requirements be met?

- A. Deploy the database on multiple Amazon EC2 instances backed by Amazon EBS across multiple Availability Zones.
- B. Use Amazon RDS with a multiple Availability Zone option.
- C. Use RDS with a single Available Zone option and schedule periodic database snapshots.
- D. Use Amazon DynamoDB.

Correct Answer: B

Reference: <https://aws.amazon.com/rds/>

QUESTION 3

A Solutions Architect is designing a microservices-based application using Amazon ECS. The application includes a WebSocket component, and the traffic needs to be distributed between microservices based on the URL. Which service should the Architect choose to distribute the workload?

- A. ELB Classic Load Balancer
- B. Amazon Route 53 DNS
- C. ELB Application Load Balancer
- D. Amazon CloudFront

Correct Answer: C

Reference: <https://docs.aws.amazon.com/aws-technical-content/latest/microservices-on-aws/microservices-on->

QUESTION 4

A Solutions Architect is developing a new web application on AWS. The services must scale to support an increasing load. The Architect wants to focus on software development and deploying new features rather than provisioning or managing servers.

Which AWS service is appropriate?

- A. Auto Scaling
- B. Elastic Beanstalk
- C. EC2 Container Service
- D. CloudFormation

Correct Answer: B

Reference: <https://aws.amazon.com/elasticbeanstalk/>

QUESTION 5

Your customer wishes to deploy an enterprise application to AWS which will consist of several web servers, several application servers and a small (50GB) Oracle database information is stored, both in the database and the file systems of

the various servers. The backup system must support database recovery whole server and whole disk restores, and individual file restores with a recovery time of no more than two hours. They have chosen to use RDS Oracle as the database.

Which backup architecture will meet these requirements?

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

Correct Answer: A

You need to use enterprise backup software to provide file level restore. See https://d0.awsstatic.com/whitepapers/Backup_and_Recovery_Approaches_Using_AWS.pdf Page 18:

If your existing backup software does not natively support the AWS cloud, you can use AWS storage gateway products.

AWS Storage Gateway is a virtual appliance that provides seamless and secure integration between your data center and the AWS storage infrastructure.

QUESTION 6

A company's Amazon RDS MySQL DB instance may be rebooted for maintenance and to apply patches. This database is critical and potential user disruption must be minimized. What should the Solution Architect do in this scenario?

- A. Set up an RDS MySQL cluster
- B. Create an RDS MySQL Read Replica.
- C. Set RDS MySQL to Multi-AZ.
- D. Create an Amazon EC2 instance MySQL cluster.

Correct Answer: D

QUESTION 7

A Solution Architect is designing a disaster recovery solution for a 5 TB Amazon Redshift cluster. The recovery site must be at least 500 miles (805 kilometers) from the live site. How should the Architect meet these requirements?

- A. Use AWS CloudFormation to deploy the cluster in a second region.
- B. Take a snapshot of the cluster and copy it to another Availability Zone.
- C. Modify the Redshift cluster to span two regions.
- D. Enable cross-region snapshots to a different region.

Correct Answer: D

QUESTION 8

An application stack includes an Elastic Load Balancer in a public subnet, a fleet of Amazon EC2 instances in an Auto Scaling group, and an Amazon RDS MySQL cluster. Users connect to the application from the Internet. The application servers and database must be secure.

How should a Solutions Architect perform this task?

- A. Create a private subnet for the Amazon EC2 instances and a public subnet for the Amazon RDS cluster.
- B. Create a private subnet for the Amazon EC2 instances and a private subnet for the Amazon RDS cluster.
- C. Create a public subnet for the Amazon EC2 instances and a private subnet for the Amazon RDS cluster.
- D. Create a public subnet for the Amazon EC2 instances and a public subnet for the Amazon RDS cluster.

Correct Answer: C

QUESTION 9

An application uses an Amazon SQS queue as a transport mechanism to deliver data to a group of EC2 instances for processing. The application owner wants to add a mechanism to archive the incoming data without modifying application code on the EC2 instances.

How can this application be re-architected to archive the data without modifying the processing instances?

- A. Trigger a Lambda function by using Amazon CloudWatch Events to retrieve messages from the SQS queue and archive to Amazon S3.
- B. Use an Amazon SNS topic to fan out the data to the SQS queue in addition to a Lambda function that records the data to an S3 bucket.
- C. Set up an Amazon Kinesis Data Stream so that multiple instances can receive data. Add a separate EC2 instance that is configured to archive all data it receives.
- D. Write the data to an S3 bucket, and use an SQS queue for S3 event notifications to tell the instances where to retrieve the data.

Correct Answer: B

QUESTION 10

An organization regularly backs up their application data. The application backups are required to be stored on Amazon S3 for a certain amount of time. The backups should be accessed instantly in the event of a disaster recovery. Which of the following Amazon S3 storage classes would be the MOST cost-effective option to meet the needs of this scenario?

- A. Glacier Storage Class
- B. Standard Storage Class
- C. Standard - Infrequent Access (IA)
- D. Reduced Redundancy Class (RRS)

Correct Answer: C

Reference: <https://aws.amazon.com/s3/features/>

QUESTION 11

A company plans to deploy a new application in AWS that reads and writes information to a database. The company wants to deploy the application in two different AWS Regions in an active-active configuration. The databases need to replicate to keep information in sync.

What should be used to meet these requirements?

- A. Amazon Athena with Amazon S3 cross-region replication
-

- B. AWS Database Migration Service with change data capture
- C. Amazon DynamoDB with global tables
- D. Amazon RDS for PostgreSQL with a cross-region Read Replica

Correct Answer: D

QUESTION 12

A Solutions Architect is designing a web application that will be hosted on Amazon EC2 instances in a public subnet. The web application uses a MySQL database in a private subnet. The database should be accessible to database administrators.

Which of the following options should the Architect recommend? (Choose two.)

- A. Create a bastion host in a public subnet, and use the bastion host to connect to the database.
- B. Log in to the web servers in the public subnet to connect to the database.
- C. Perform DB maintenance after using SSH to connect to the NAT Gateway in a public subnet.
- D. Create an IPSec VPN tunnel between the customer site and the VPC, and use the VPN tunnel to connect to the database.
- E. Attach an Elastic IP address to the database.

Correct Answer: BD

QUESTION 13

A customer owns a MySQL database that is accessed by various clients who expect, at most, 100 ms latency on requests. Once a record is stored in the database, it rarely changed. Clients only access one record at a time.

Database access has been increasing exponentially due to increased client demand. The resultant load will soon exceed the capacity of the most expensive hardware available for purchase. The customer wants to migrate to AWS, and is

willing to change database systems.

Which service would alleviate the database load issue and offer virtually unlimited scalability for the future?

- A. Amazon RDS
- B. Amazon DynamoDB
- C. Amazon Redshift
- D. AWS Data Pipeline

Correct Answer: B

QUESTION 14

A company deployed a three-tier web application on Amazon EBS backed Amazon EC2 instances for the web and application tiers, and Amazon RDS for the database tier. The company is concerned about loss of data in the web and application tiers.

What is the MOST efficient way to prevent data loss?

- A. Create an Amazon EFS file system and run a shell script to copy the data
- B. Create an Amazon EBS snapshot using an Amazon CloudWatch Events rule
- C. Create an Amazon S3 snapshot policy to back up the Amazon EBS volumes
- D. Create a snapshot lifecycle policy that takes periodic snapshots of the Amazon EBS volumes

Correct Answer: B

Reference <https://aws.amazon.com/ebs/>

QUESTION 15

Two Auto Scaling applications, Application A and Application B, currently run within a shared set of subnets. A Solutions Architect wants to make sure that Application A can make requests to Application B, but Application B should be denied from making requests to Application A.

Which is the SIMPLEST solution to achieve this policy?

- A. Using security groups that Reference: the security groups of the other application
- B. Using security groups that Reference: the application server's IP addresses
- C. Using Network Access Control Lists to allow/deny traffic based on application IP addresses
- D. Migrating the applications to separate subnets from each other

Correct Answer: A

QUESTION 16

An application runs on Amazon EC2 instances in multiple Availability Zones (AZs) behind an Application Load Balancer. The load balancer is in public subnets; the EC2 instances are in private subnets and must not be accessible from the internet. The EC2 instances must call external services on the internet. If one AZ becomes unavailable, the remaining EC2 instances must still be able to call the external services.

How should these requirements be met?

- A. Create a NAT gateway attached to the VPC. Add a route to the gateway to each private subnet route table
- B. Configure an internet gateway. Add a route to the gateway to each private subnet route table.

C. Create a NAT instance in the private subnet of each AZ. Update the route tables for each private subnet to direct internet-bound traffic to the NAT instance.

D. Create a NAT gateway in each AZ. Update the route tables for each private subnet to direct internet-bound traffic to the NAT gateway.

Correct Answer: D

QUESTION 17

A company plans to use an Amazon VPC to deploy a web application consisting of an elastic load balancer, a fleet of web and application servers, and an Amazon RDS MySQL database that should not be accessible from the Internet. The proposed design must be highly available and distributed over two Availability Zones.

What would be the MOST appropriate VPC design for this specific use case?

- A. Two public subnets for the elastic load balancer, two public subnets for the web servers, and two public subnets for Amazon RDS.
- B. One public subnet for the elastic load balancer, two private subnets for the web servers, and two private subnets for Amazon RDS.
- C. One public subnet for the elastic load balancer, one public subnet for the web servers, and one private subnet for the database.
- D. Two public subnets for the elastic load balancer, two private subnets for the web servers, and two private subnets for RDS.

Correct Answer: B

QUESTION 18

A company is migrating on-premises databases to AWS. The company's backend application produces a large amount of database queries for reporting purposes, and the company wants to offload some of those reads to Read Replica, allowing the primary database to continue performing efficiently.

Which AWS database platforms will accomplish this? (Select TWO.)

- A. Amazon RDS for Oracle
- B. Amazon RDS for PostgreSQL
- C. Amazon RDS for MariaDB
- D. Amazon DynamoDB
- E. Amazon RDS for Microsoft SQL Server

Correct Answer: BC

Reference: <http://jayendrapatil.com/aws-rds-replication-multi-az-read-replica/>

QUESTION 19

A Solutions Architect is designing a Lambda function that calls an API to list all running Amazon RDS instances. How should the request be authorized?

- A. Create an IAM access and secret key, and store it in the Lambda function.
- B. Create an IAM role to the Lambda function with permissions to list all Amazon RDS instances.
- C. Create an IAM role to Amazon RDS with permissions to list all Amazon RDS instances.
- D. Create an IAM access and secret key, and store it in an encrypted RDS database.

Correct Answer: C

QUESTION 20

A company is migrating its data center to AWS. As part of this migration, there is a three-tier web application that has strict data-at-rest encryption requirements. The customer deploys this application on Amazon EC2 using Amazon EBS, and now must provide encryption at-rest.

How can this requirement be met without changing the application?

- A. Use AWS Key Management Service and move the encrypted data to Amazon S3.
- B. Use an application-specific encryption API with AWS server-side encryption.
- C. Use encrypted EBS storage volumes with AWS-managed keys.
- D. Use third-party tools to encrypt the EBS data volumes with Key Management Service Bring Your Own Keys.

Correct Answer: C

QUESTION 21

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

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

Correct Answer: A

A single-node Memcached ElastiCache cluster failure is nothing but a total failure. (Even though AWS will automatically recover the failed node, there are no other nodes in the cluster) <http://docs.aws.amazon.com/AmazonElastiCache/latest/UserGuide/BestPractices.html> Mitigating Node Failures To mitigate the impact of a node failure, spread your cached data over more nodes. Because Memcached does not support replication, a node failure will always result in some data loss from your cluster. When you create your Memcached cluster you can create it with 1 to 20 nodes, or more by special request. Partitioning your data across a greater number of nodes means you'll lose less data if a node fails. For example, if you partition your data across 10 nodes, any single node stores approximately 10% of your cached data. In this case, a node failure loses approximately 10% of your cache which needs to be replaced when a replacement node is created and provisioned. Mitigating Availability Zone Failures

To mitigate the impact of an availability zone failure, locate your nodes in as many availability zones as possible. In the unlikely event of an AZ failure, you will lose only the data cached in that AZ, not the data cached in the other AZs.

QUESTION 22

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

and Keep costs to a minimum.

What AWS architecture would you recommend?

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

Correct Answer: A

In question we have keywords `scalable` and company wants to `move systems` to AWS, which is best suited for Auto-scaling group. <https://aws.amazon.com/blogs/security/writing-iam-policies-grant-access-to-user-specific-folders-in-an-amazon-s3-bucket/>

QUESTION 23

A Solutions Architect is creating an application running in an Amazon VPC that needs to access AWS Systems Manager Parameter Store. Network security rules prohibit any route table entry with a 0.0.0.0/0 destination.

What infrastructure addition will allow access to the AWS service while meeting the requirements?

- A. VPC peering

- B. NAT instance
- C. NAT gateway
- D. AWS PrivateLink

Correct Answer: A

QUESTION 24

A Solutions Architect is designing an application on AWS that uses persistent block storage. Data must be encrypted at rest.

Which solution meets the requirement?

- A. Enable SSL on Amazon EC2 instances.
- B. Encrypt Amazon EBS volumes on Amazon EC2 instances.
- C. Enable server-side encryption on Amazon S3.
- D. Encrypt Amazon EC2 Instance Storage.

Correct Answer: B

Reference: <https://aws.amazon.com/blogs/aws/protect-your-data-with-new-ebs-encryption/>

QUESTION 25

A Solutions Architect must design a solution that encrypts data in Amazon S3. Corporate policy mandates encryption keys be generated and managed on premises. Which solution should the Architect use to meet the security requirements?

- A. AWS CloudHSM
- B. SSE-KMS: Server-side encryption with AWS KMS managed keys
- C. SSE-S3: Server-side encryption with Amazon-managed master key
- D. SSE-C: Server-side encryption with customer-provided encryption keys

Correct Answer: B

QUESTION 26

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

A. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.

B. Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. EBS volumes to host videos and EBS snapshots to incrementally backup original files after a few days. CloudFront to serve HLS transcoded videos from EC2.

C. Elastic Transcoder to transcode original high-resolution MP4 videos to HLS. S3 to host videos with Lifecycle Management to archive original files to Glacier after a few days. CloudFront to serve HLS transcoded videos from S3.

D. A video transcoding pipeline running on EC2 using SQS to distribute tasks and Auto Scaling to adjust the number of nodes depending on the length of the queue. S3 to host videos with Lifecycle Management to archive all files to Glacier after a few days. CloudFront to serve HLS transcoded videos from Glacier.

Correct Answer: C

QUESTION 27

Amazon EC2 provides virtual computing environments known as _____.

A. instances

B. volumes

C. microsystems

D. servers

Correct Answer: A

Amazon EC2 provides virtual computing environments known as instances. When you launch an instance, the instance type that you specify determines the hardware of the host computer used for your instance. Each instance type offers different compute, memory, and storage capabilities and are grouped in instance families based on these capabilities. Select an instance type based on the requirements of the application or software that you plan to run on your instance. <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html>

QUESTION 28

Which tool analyzes account resources and provides a detailed inventory of changes over time?

A. AWS Config

B. AWS CloudFormation

C. Amazon CloudWatch

D. AWS Service Catalog

Correct Answer: A

Reference: <https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>

QUESTION 29

Your company currently has a 2-tier web application running in an on-premises data center. You have experienced several infrastructure failures in the past two months resulting in significant financial losses. Your CIO is strongly agreeing to

move the application to AWS. While working on achieving buy-in from the other company executives, he asks you to develop a disaster recovery plan to help improve Business continuity in the short term. He specifies a target Recovery Time

Objective (RTO) of 4 hours and a Recovery Point Objective (RPO) of 1 hour or less. He also asks you to implement the solution within 2 weeks. Your database is 200GB in size and you have a 20Mbps Internet connection.

How would you do this while minimizing costs?

- A. Create an EBS backed private AMI which includes a fresh install of your application. Develop a CloudFormation template which includes your AMI and the required EC2, AutoScaling, and ELB resources to support deploying the application across Multiple- Availability-Zones. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- B. Deploy your application on EC2 instances within an Auto Scaling group across multiple availability zones. Asynchronously replicate transactions from your on-premises database to a database instance in AWS across a secure VPN connection.
- C. Create an EBS backed private AMI which includes a fresh install of your application. Setup a script in your data center to backup the local database every 1 hour and to encrypt and copy the resulting file to an S3 bucket using multi-part upload.
- D. Install your application on a compute-optimized EC2 instance capable of supporting the application's average load. Synchronously replicate transactions from your on-premises database to a database instance in AWS across a secure Direct Connect connection.

Correct Answer: A

Overview of Creating Amazon EBS-Backed AMIs First, launch an instance from an AMI that's similar to the AMI that you'd like to create. You can connect to your instance and customize it. When the instance is configured correctly, ensure data integrity by stopping the instance before you create an AMI, then create the image. When you create an Amazon EBS-backed AMI, we automatically register it for you. Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can tell Amazon EC2 not to power down and reboot the instance. Some file systems, such as XFS, can freeze and unfreeze activity, making it safe to create the image without rebooting the instance. During the AMI-creation process, Amazon EC2 creates snapshots of your instance's root volume and any other EBS volumes attached to your instance. If any volumes attached to the instance are encrypted, the new AMI only launches successfully on instances that support Amazon EBS encryption. For more information, see Amazon EBS Encryption. Depending on the size of the volumes, it can take several minutes for the AMI-creation process to complete (sometimes up to 24 hours). You may find it more efficient to create snapshots of your volumes prior to creating your AMI. This way, only small, incremental snapshots need to be created when the AMI is created, and the process completes more quickly (the total time for snapshot creation remains the same). For more information, see Creating an Amazon EBS Snapshot. After the process completes, you have a new AMI and snapshot created from the root volume of the instance. When you launch an instance using the new AMI, we create a new EBS volume for its root volume using the snapshot. Both the AMI and the snapshot incur charges to your account until you delete them. For more information, see Deregistering Your AMI. If you add instance-store volumes or EBS volumes to your instance in addition to the root device volume, the block device mapping for the new AMI contains information for these volumes, and the block device mappings for instances that you launch from the new AMI automatically contain information for these volumes. The instance-store volumes specified in the block device mapping for the new instance

are new and don't contain any data from the instance store volumes of the instance you used to create the AMI. The data on EBS volumes persists. For more information, see Block Device Mapping.

QUESTION 30

A company has an application that uses Amazon CloudFront for content that is hosted on an Amazon S3 bucket. After an unexpected refresh, the users are still seeing old content. Which step should the Solutions Architect take to ensure that new content is displayed?

- A. Perform a cache refresh on the CloudFront distribution that is serving the content.
- B. Perform an invalidation on the CloudFront distribution that is serving the content.
- C. Create a new cache behavior path with the updated content.
- D. Change the TTL value for removing the old objects.

Correct Answer: D

QUESTION 31

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

- A. Use RDS Multi-AZ with two tables, one for "Active calls" and one for "Terminated calls". In this way the "Active calls" table is always small and effective to access.
- B. Use DynamoDB with a "Calls" table and a Global Secondary Index on a "IsActive" attribute that is present for active calls only. In this way the Global Secondary index is sparse and more effective.
- C. Use DynamoDB with a "Calls" table and a Global secondary index on a "State" attribute that can equal to "active" or "terminated" in this way the Global Secondary index can be used for all items in the table.
- D. Use RDS Multi-AZ with a "CALLS" table and an Indexed "STATE" field that can be equal to "ACTIVE" or "TERMINATED". In this way the SQL query is optimized by the use of the Index.

Correct Answer: B

Q:

Can a global secondary index key be defined on non-unique attributes?

Yes. Unlike the primary key on a table, a GSI index does not require the indexed attributes to be unique.

Q:

Are GSI key attributes required in all items of a DynamoDB table?

No. GSIs are sparse indexes. Unlike the requirement of having a primary key, an item in a DynamoDB table does not have to contain any of the GSI keys. If a GSI key has both hash and range elements, and a table item omits either of them,

then that item will not be indexed by the corresponding GSI. In such cases, a GSI can be very useful in efficiently locating items that have an uncommon attribute.

QUESTION 32

A Solutions Architect is designing a new application that needs to access data in a different AWS account located within the same region. The data must not be accessed over the Internet. Which solution will meet these requirements with the LOWEST cost?

- A. Add rules to the security groups in each account.
- B. Establish a VPC Peering connection between accounts.
- C. Configure Direct Connect in each account.
- D. Add a NAT Gateway to the data account.

Correct Answer: B

QUESTION 33

An application running in a private subnet accesses an Amazon DynamoDB table. There is a security requirement that the data never leave the AWS network. How should this requirement be met?

- A. Configure a network ACL on DynamoDB to limit traffic to the private subnet
- B. Enable DynamoDB encryption at rest using an AWS KMS key
- C. Add a NAT gateway and configure the route table on the private subnet
- D. Create a VPC endpoint for DynamoDB and configure the endpoint policy

Correct Answer: C

reference https://docs.aws.amazon.com/vpc/latest/userguide/VPC_Scenario2.html

QUESTION 34

A company needs to use AWS resources to expand capacity for a website hosted in an on-premises data center. The AWS resources will include load balancers, Auto Scaling, and Amazon EC2 instances that will access an on-premises database. Network connectivity has been established, but no traffic is going to the AWS environment.

How should Amazon Route 53 be configured to distribute load to the AWS environment? (Select TWO.)

- A. Set up a weighted routing policy, distributing the workload between the load balancer and the on-premises environment.

- B. Set up an A record to point the DNS name to the IP address of the load balancer.
- C. Create multiple A records for the EC2 instances.
- D. Set up a geolocation routing policy to distribute the workload between the load balancer and the on-premises environment.
- E. Set up a routing policy for failover using the on-premises environment as primary and the load balancer as secondary.

Correct Answer: AB

QUESTION 35

A Solutions Architect is designing a solution for a dynamic website, "example.com," that is deployed in two regions: Tokyo, Japan and Sydney, Australia. The Architect wants to ensure that users located in Australia are directed to the website deployed in the Sydney region and users located in Japan are redirected to the website in the Tokyo region when they browse to "example.com".

Which service should the Architect use to achieve this goal with the LEAST administrative effort?

- A. Amazon CloudFront with geolocation routing
- B. Amazon Route 53
- C. Application Load Balancer
- D. Network Load Balancer deployed across multiple regions

Correct Answer: B

Reference: <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html>

QUESTION 36

When designing an Amazon SQS message-processing solution, messages in the queue must be processed before the maximum retention time has elapsed. Which actions will meet this requirement? (Choose two.)

- A. Use AWS STS to process the messages
- B. Use Amazon EBS-optimized Amazon EC2 instances to process the messages
- C. Use Amazon EC2 instances in an Auto Scaling group with scaling triggered based on the queue length
- D. Increase the SQS queue attribute for the message retention period
- E. Convert the SQS queue to a first-in first-out (FIFO) queue

Correct Answer: BD

QUESTION 37

A Solutions Architect is designing a solution that includes a managed VPN connection. To monitor whether the VPN connection is up or down, the Architect should use:

- A. an external service to ping the VPN endpoint from outside the VPC.
- B. AWS CloudTrail to monitor the endpoint.
- C. the CloudWatch TunnelState Metric.
- D. an AWS Lambda function that parses the VPN connection logs.

Correct Answer: C

Reference: <https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/monitoring-cloudwatch-vpn.html>

QUESTION 38

An organization runs an online media site, hosted on-premises. An employee posted a product review that contained videos and pictures. The review went viral and the organization needs to handle the resulting spike in website traffic. What action would provide an immediate solution?

- A. Redesign the website to use Amazon API Gateway, and use AWS Lambda to deliver content.
- B. Add server instances using Amazon EC2 and use Amazon Route 53 with a failover routing policy.
- C. Serve the images and videos via an Amazon CloudFront distribution created using the news site as the origin.
- D. Use Amazon ElasticCache for Redis for caching and reducing the load requests from the origin.

Correct Answer: C

QUESTION 39

You are designing the network infrastructure for an application server in Amazon VPC. Users will access all the application instances from the Internet as well as from an on-premises network. The on-premises network is connected to your

VPC over an AWS Direct Connect link.

How would you design routing to meet the above requirements?

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

Correct Answer: B

QUESTION 40

A Solutions Architect must review an application deployed on EC2 instances that currently stores multiple 5-GB files on attached instance store volumes. The company recently experienced a significant data loss after stopping and starting their instances and wants to prevent the data loss from happening again. The solution should minimize performance impact and the number of code changes required.

What should the Solutions Architect recommend?

- A. Store the application data in Amazon S3
- B. Store the application data in an EBS volume
- C. Store the application data in Amazon ElastiCache
- D. Store the application data in Amazon DynamoDB

Correct Answer: B

Reference: <http://jayendrapatil.com/tag/i2/>