Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1
You run a clustered NoSQL database on AWS EC2 using AWS EBS. You need to reduce latency for database response times. Performance is the most important concern, not availability. You did not perform the initial setup, someone without much AWS knowledge did, so you are not sure if they configured everything optimally. Which of the following is NOT likely to be an issue contributing to increased latency?

A. The EC2 instances are not EBS Optimized.
B. The database and requesting system are both in the wrong Availability Zone.
C. The EBS Volumes are not using PIOPS.
D. The database is not running in a placement grou

Answer: B

Explanation:
For the highest possible performance, all instances in a clustered database like this one should be in a single Availability Zone in a placement group, using EBS optimized instances, and using PIOPS SSD EBS Volumes. The particular Availability Zone the system is running in should not be important, as long as it is the same as the requesting resources.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 2
When thinking of DynamoDB, what are true of Global Secondary Key properties?

A. The partition key and sort key can be different from the table.
B. Only the partition key can be different from the table.
C. Either the partition key or the sort key can be different from the table, but not both.
D. Only the sort key can be different from the tabl

Answer: A

Explanation:
Global secondary index — an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.
Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html

NEW QUESTION 3
You are designing a service that aggregates clickstream data in batch and delivers reports to subscribers via email only once per week. Data is extremely spikey, geographically distributed, high-scale, and unpredictable. How should you design this system?

A. Use a large RedShift cluster to perform the analysis, and a fileet of Lambdas to perform record inserts into the RedShift table
B. Lambda will scale rapidly enough for the traffic spikes.
C. Use a CloudFront distribution with access log delivery to S3. Clicks should be recorded as querystring GETs to the distributio
D. Reports are built and sent by periodically running EMR jobs over the access logs in S3.
E. Use API Gateway invoking Lambdas which PutRecords into Kinesis, and EMR running Spark performing GetRecords on Kinesis to scale with spike
F. Spark on EMR outputs the analysis to S3, which are sent out via email.
G. Use AWS Elasticsearch service and EC2 Auto Scaling group
H. The Autoscaling groups scale based on click throughput and stream into the Elasticsearch domain, which is also scalabl
I. Use Kibana to generate reports periodically.

Answer: B

Explanation:
Because you only need to batch analyze, anything using streaming is a waste of money. CloudFront is a Gigabit-Scale HTTP(S) global request distribution service, so it can handle scale, geo-spread, spikes, and unpredictability. The Access Logs will contain the GET data and work just fine for batch analysis and email using EMR.
Can I use Amazon CloudFront if I expect usage peaks higher than 10 Gbps or 15,000 RPS? Yes. Complete our request for higher limits here, and we will add more capacity to your account within two business days.
Reference: https://aws.amazon.com/Cloudfront/faqs/

NEW QUESTION 4
You need to create a simple, holistic check for your system's general availablity and uptime. Your system presents itself as an HTTP-speaking API. What is the most simple tool on AWS to achieve this with?

A. Route53 Health Checks
B. CloudWatch Health Checks
C. AWS ELB Health Checks
D. EC2 Health Checks

Answer: A

Explanation:
You can create a health check that will run into perpetuity using Route53, in one API call, which will ping your service via HTTP every 10 or 30 seconds.
Amazon Route 53 must be able to establish a TCP connection with the endpoint within four seconds. In addition, the endpoint must respond with an HTTP status code of 200 or greater and less than 400 within two seconds after connecting.
Reference:
http://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-determining-health-of-endpoint s.html

NEW QUESTION 5
Which EBS volume type is best for high performance NoSQL cluster deployments?

A. io1
B. gp1
C. standard
D. gp2

Answer: A

Explanation:
io1 volumes, or Provisioned IOPS (PIOPS) SSDs, are best for: Critical business applications that require sustained IOPS performance, or more than 10,000 IOPS or 160 MiB/s of throughput per volume, like large database workloads, such as MongoDB.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVo|umeTypes.htm|

NEW QUESTION 6
Which status represents a failure state in AWS CIoudFormation?

A. <code>UPDATE_COMPLETE_CLEANUP_IN_PROGRESS</code>
B. <code>DELETE_COMPLETE_WITH_ARTIFACTS</code>
C. <code>ROLLBACK_IN_PROGRESS</code>
D. <code>ROLLBACK_FAILED</code>

Answer: C

Explanation:
ROLLBACK_IN_PROGRESS means an UpdateStack operation failed and the stack is in the process of trying to return to the valid, pre-update state.
UPDATE_COMPLETE_CLEANUP_IN_PROGRESS means an update was successful, and CIoudFormation is deleting any replaced, no longer used resources.
ROLLBACK_FA|LED is not a CloudFormation state (but UPDATE_ROLLBACK_FAILED is). DELETE_COMPLETE_W|TH_ART|FACTS does not exist at all.
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/using-cfn-updating-stacks.html

NEW QUESTION 7
Which deployment method, when using AWS Auto Scaling Groups and Auto Scaling Launch Configurations, enables the shortest time to live for indMdual sewers?

A. Pre-baking AMIs with all code and configuration on deploys.
B. Using a Dockerfile bootstrap on instance launch.
C. Using UserData bootstrapping scripts.
D. Using AWS EC2 Run Commands to dynamically SSH into fileet

Answer: A

Explanation:
Note that the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fileet of applications with several build tools and dependencies can be a challenging task during rollouts. Furthermore, your deployment service should be designed to do faster rollouts to take advantage of Auto Scaling. Prebaking is a process of embedding a significant portion of your application artifacts within your base AMI. During the deployment process you can customize application installations by using EC2 instance artifacts such as instance tags, instance metadata, and Auto Scaling groups.
Reference: https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf

NEW QUESTION 8
Which of the following are not valid sources for OpsWorks custom cookbook repositories?

A. HTTP(S)
B. Git
C. AWS EBS
D. Subversion

Answer: C

Explanation:
Linux stacks can install custom cookbooks from any of the following repository types: HTTP or Amazon S3 archives. They can be either public or private, but Amazon S3 is typically the preferred option for a private archive. Git and Subversion repositories provide source control and the ability to have multiple versions.
Reference:
http://docs.aws.amazon.com/opsworks/latest/userguide/workingcookbook-instaIIingcustom-enable.html

NEW QUESTION 9
You are building a deployment system on AWS. You will deploy new code by bootstrapping instances in a private subnet in a VPC at runtime using UserData scripts pointing to an S3 zip file object, where your code is stored. An ELB in a public subnet has network interfaces and connectMty to the instances. Requests from users of the system are routed to the ELB via a Route53 A Record Alias. You do not use any VPC endpoints. Which is a risk of using this approach?

A. Route53 Alias records do not always update dynamically with ELB network changes after deploys.
B. If the NAT routing for the private subnet fails, deployments fail.
C. Kernel changes to the base AMI may render the code inoperable.
D. The instances cannot be in a private subnet if the ELB is in a public on

Answer: B

Explanation:
Since you are not using VPC endpoints, outbound requests for the code sitting in S3 are routed though the NAT for the VPC's private subnets. If this networking fails, runtime bootstrapping through code

download will fail due to network unavailability and lack of access to the Internet, and thus Amazon S3. Reference:
http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html

NEW QUESTION 10
You need to grant a vendor access to your AWS account. They need to be able to read protected messages in a private S3 bucket at their leisure. They also use AWS. What is the best way to accomplish this?

A. Create an IAM User with API Access Key
B. Grant the User permissions to access the bucke
C. Give the vendor the AWS Access Key ID and AWS Secret Access Key for the User.
D. Create an EC2 Instance Profile on your accoun
E. Grant the associated IAM role full access to the bucke
F. Start an EC2 instance with this Profile and give SSH access to the instance to the vendor.
G. Create a cross-account IAM Role with permission to access the bucket, and grant permission to use the Role to the vendor AWS account.
H. Generate a signed S3 PUT URL and a signed S3 PUT URL, both with wildcard values and 2 year duration
I. Pass the URLs to the vendor.

Answer: C

Explanation:
When third parties require access to your organization's AWS resources, you can use roles to delegate access to them. For example, a third party might provide a service for managing your AWS resources. With IAM roles, you can grant these third parties access to your AWS resources without sharing your AWS security credentials. Instead, the third party can access your AWS resources by assuming a role that you create in your AWS account.
Reference:
http://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_common-scenarios_third-party.html

NEW QUESTION 10
Your serverless architecture using AWS API Gateway, AWS Lambda, and AWS DynamoDB experienced a large increase in traffic to a sustained 400 requests per second, and dramatically increased in failure rates. Your requests, during normal operation, last 500 milliseconds on average. Your DynamoDB table did not exceed 50% of provisioned throughput, and Table primary keys are designed correctly. What is the most likely issue?

A. Your API Gateway deployment is throttling your requests.
B. Your AWS API Gateway Deployment is bottlenecking on request (de)serialization.
C. You did not request a limit increase on concurrent Lambda function executions.
D. You used Consistent Read requests on DynamoDB and are experiencing semaphore loc

Answer: C

Explanation:
AWS API Gateway by default throttles at 500 requests per second steady-state, and 1000 requests per second at spike. Lambda, by default, throttles at 100 concurrent requests for safety. At 500 milliseconds (half of a second) per request, you can expect to support 200 requests per second at 100 concurrency. This is less than the 400 requests per second your system now requires. Make a limit increase request via the AWS Support Console.
AWS Lambda: Concurrent requests safety throttle per account -> 100
Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.htm|#|imits_|ambda

NEW QUESTION 15
For AWS CloudFormation, which stack state refuses UpdateStack calls?

A. <code>UPDATE_ROLLBACK_FAILED</code>
B. <code>UPDATE_ROLLBACK_COMPLETE</code>
C. <code>UPDATE_CONIPLETE</code>
D. <code>CREATE_COMPLETE</code>

Answer: A

Explanation:
When a stack is in the UPDATE_ROLLBACK_FA|LED state, you can continue rolling it back to return it to a working state (to UPDATE_ROLLBACK_COMPLETE). You cannot update a stack that is in the UPDATE_ROLLBACK_FA|LED state. However, if you can continue to roll it back, you can return the stack to its original settings and try to update it again.
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/using-cfn-updating-stacks-continueu pdateroIlback.html

NEW QUESTION 19
You are experiencing performance issues writing to a DynamoDB table. Your system tracks high scores for video games on a marketplace. Your most popular game experiences all of the performance issues. What is the most likely problem?

A. DynamoDB's vector clock is out of sync, because of the rapid growth in request for the most popular game.
B. You selected the Game ID or equivalent identifier as the primary partition key for the table.
C. Users of the most popular video game each perform more read and write requests than average.
D. You did not provision enough read or write throughput to the tabl

Answer: B

Explanation:
The primary key selection dramatically affects performance consistency when reading or writing to DynamoDB. By selecting a key that is tied to the identity of the game, you forced DynamoDB to create a hotspot in the table partitions, and over-request against the primary key partition for the popular game. When it stores data, DynamoDB dMdes a table's items into multiple partitions, and distributes the data primarily based upon the partition key value. The provisioned throughput associated with a table is also dMded evenly among the partitions, with no sharing of provisioned throughput across partitions. Reference:
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#Guideli nesForTables.UniformWorkload

NEW QUESTION 20
When thinking of AWS OpsWorks, which of the following is not an instance type you can allocate in a stack layer?

A. 24/7 instances
B. Spot instances
C. Time-based instances
D. Load-based instances

Answer: B

Explanation:
AWS OpsWorks supports the following instance types, which are characterized by how they are started and stopped. 24/7 instances are started manually and run until you stop them.Time-based instances are run by AWS OpsWorks on a specified daily and weekly schedule. They allow your stack to automatically adjust the number of instances to accommodate predictable usage patterns. Load-based instances are automatically started and stopped by AWS OpsWorks, based on specified load metrics, such as CPU utilization. They allow your stack to automatically adjust the number of instances to accommodate variations in incoming traffic. Load-based instances are available only for Linux-based stacks. Reference: http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html

NEW QUESTION 25
Which of these is not a CloudFormation Helper Script?

A. cfn-signal
B. cfn-hup
C. cfn-request
D. cfn-get-metadata

Answer: C

Explanation:
This is the complete list of CloudFormation Helper Scripts: cfn-init, cfn-signal, cfn-get-metadata, cfn-hup Reference:
http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scripts-reference.html

NEW QUESTION 26
You need to replicate API calls across two systems in real time. What tool should you use as a buffer and transport mechanism for API call events?

A. AWS SQS
B. AWS Lambda
C. AWS Kinesis
D. AWS SNS

Answer: C

Explanation:
AWS Kinesis is an event stream service. Streams can act as buffers and transport across systems for in-order programmatic events, making it ideal for replicating API calls across systems.
A typical Amazon Kinesis Streams application reads data from an Amazon Kinesis stream as data records. These applications can use the Amazon Kinesis Client Library, and they can run on Amazon EC2 instances. The processed records can be sent to dashboards, used to generate alerts, dynamically change pricing and advertising strategies, or send data to a variety of other AWS services. For information about Streams features and pricing, see Amazon Kinesis Streams.
Reference: http://docs.aws.amazon.com/kinesis/latest/dev/introduction.html

NEW QUESTION 27
Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as possible. How should you do this?

A. Create a global AWS CloudTrail Trai
B. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.
C. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API call
D. Subscribe the CTO to an email type delivery on this SNS Topic.
E. Use AWS IAM credential reports to deliver a CSV of all uses of IAM User Tokens over time to the CTO.
F. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB tabl
G. Generate reports based on the contents of this table.

Answer: A

Explanation:
This is the ideal use case for AWS CloudTrail.
CloudTrail provides visibility into user actMty by recording API calls made on your account. CloudTrail records important information about each API call, including the name of the API, the identity of the caller, the time of the API call, the request parameters, and the response elements returned by the AWS service. This information helps you to track changes made to your AWS resources and to troubleshoot operational issues. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards. Reference: https://aws.amazon.com/Cloudtrail/faqs/

NEW QUESTION 29
What is the order of most-to-least rapidly-scaling (fastest to scale first)?
(A) EC2 + ELB + Auto Scaling (B) Lambda (C) RDS

A. B, A, C
B. C, B, A
C. C, A, B

D. A, C, B

Answer: A

Explanation:
Lambda is designed to scale instantly. EC2 + ELB + Auto Scaling require single-digit minutes to scale out. RDS will take atleast 15 minutes, and will apply OS patches or any other updates when applied. Reference: https://aws.amazon.com/|ambda/faqs/

NEW QUESTION 31
What is required to achieve gigabit network throughput on EC2? You already selected cluster-compute, 10GB instances with enhanced networking, and your workload is already network-bound, but you are not seeing 10 gigabit speeds.

A. Enable biplex networking on your servers, so packets are non-blocking in both directions and there's no switching overhead.
B. Ensure the instances are in different VPCs so you don't saturate the Internet Gateway on any one VPC.
C. Select PIOPS for your drives and mount several, so you can provision sufficient disk throughput.
D. Use a placement group for your instances so the instances are physically near each other in the same Availability Zone.

Answer: D

Explanation:
You are not guaranteed 10gigabit performance, except within a placement group.
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 Gbps network. Placement groups are recommended for applications that benefit from low network latency, high network throughput, or both. Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

NEW QUESTION 34
From a compliance and security perspective, which of these statements is true?

A. You do not ever need to rotate access keys for AWS IAM Users.
B. You do not ever need to rotate access keys for AWS IAM Roles, nor AWS IAM Users.
C. None of the other statements are true.
D. You do not ever need to rotate access keys for AWS IAM Role

Answer: D

Explanation:
IAM Role Access Keys are auto-rotated by AWS on your behalf; you do not need to rotate them.
The application is granted the permissions for the actions and resources that you've defined for the role through the security credentials associated with the role. These security credentials are temporary and we
rotate them automatically. We make new credentials available at least five minutes prior to the expiration of the old credentials.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

NEW QUESTION 35
Which of these is not a reason a Mu|ti-AZ RDS instance will failover?

A. An Availability Zone outage
B. A manual failover of the DB instance was initiated using Reboot with failover
C. To autoscale to a higher instance class
D. The primary DB instance fails

Answer: C

Explanation:
The primary DB instance switches over automatically to the standby replica if any of the > following conditions occur: An Availability Zone outage, the primary DB instance fails, the DB instance's server type is changed, the operating system of the DB instance is, undergoing software patching, a manual failover
of the DB instance was initiated using Reboot with failover
Reference: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Concepts.Mu|tiAZ.htmI

NEW QUESTION 39
Which of the following tools does not directly support AWS OpsWorks, for monitoring your stacks?

A. AWS Config
B. Amazon CIoudWatch NIetrics
C. AWS CloudTraiI
D. Amazon CIoudWatch Logs

Answer: A

Explanation:
You can monitor your stacks in the following ways: AWS OpsWorks uses Amazon CIoudWatch to provide thirteen custom metrics with detailed monitoring for each instance in the stack; AWS OpsWorks integrates with AWS CIoudTrail to log every AWS OpsWorks API call and store the data in an Amazon S3 bucket; You can use Amazon CIoudWatch Logs to monitor your stack's system, application, and custom logs. Reference:
http://docs.aws.amazon.com/opsworks/latest/userguide/monitoring.htmI

NEW QUESTION 40
What is true of the way that encryption works with EBS?

A. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
B. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot creates an encrypted volume when specified / requested.
C. Snapshotting an encrypted volume makes an encrypted snapshot; restoring an encrypted snapshot always creates an encrypted volume.
D. Snapshotting an encrypted volume makes an encrypted snapshot when specified / requested; restoring an encrypted snapshot always creates an encrypted volume.

Answer: C

Explanation:
Snapshots that are taken from encrypted volumes are automatically encrypted. Volumes that are created from encrypted snapshots are also automatically encrypted. Your encrypted volumes and any associated snapshots always remain protected. For more information, see Amazon EBS Encryption.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.htmI

NEW QUESTION 43
You are designing a system which needs, at minimum, 8 m4.large instances operating to service traffic. When designing a system for high availability in the us-east-1 region, which has 6 Availability Zones, you company needs to be able to handle death of a full availability zone. How should you distribute the servers, to save as much cost as possible, assuming all of the EC2 nodes are properly linked to an ELB? Your VPC account can utilize us-east-1's AZ's a through f, inclusive.

A. 3 servers in each of AZ's a through d, inclusive.
B. 8 servers in each of AZ's a and b.
C. 2 servers in each of AZ's a through e, inclusive.
D. 4 servers in each of AZ's a through c, inclusiv

Answer: C

Explanation:
You need to design for N+1 redundancy on Availability Zones. ZONE_COUNT = (REQUIRED_INSTANCES / INSTANCE_COUNT_PER_ZONE) + 1. To minimize cost, spread the instances across as many possible zones as you can. By using a though e, you are allocating 5 zones. Using 2 instances, you have 10 total instances. If a single zone fails, you have 4 zones left, with 2 instances each, for a total of 8 instances. By spreading out as much as possible, you have increased cost by only 25% and significantly de-risked an availability zone failure.
Reference:
http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.htmI#concepts- regions-availability-zones

NEW QUESTION 47
You need to create a Route53 record automatically in CIoudFormation when not running in production during all launches of a Template. How should you implement this?

A. Use a <code>Parameter</code> for <code>environment</code>, and add a <code>Condition</code> on the Route53 <code>Resource</code> in the template to create the record only when<code>environment</code> is not <code>production</code>.
B. Create two templates, one with the Route53 record value and one with a null value for the recor
C. Use the one without it when deploying to production.
D. Use a <code>Parameter</code> for <code>environment</code>, and add a <code>Condition</code> on the Route53 <code>Resource</code> in the template to create the record with a null string when<code>environment</code> is <code>production</code>.
E. Create two templates, one with the Route53 record and one without i
F. Use the one without it when deploying to production.

Answer: A

Explanation:
The best way to do this is with one template, and a Condition on the resource. Route53 does not allow null strings for records.
Reference:
http://docs.aws.amazon.com/AWSCIoudFormation/latest/UserGuide/conditions-section-structure.html

NEW QUESTION 48
Your system uses a multi-master, multi-region DynamoDB configuration spanning two regions to achieve high availablity. For the first time since launching your system, one of the AWS Regions in which you operate over went down for 3 hours, and the failover worked correctly. However, after recovery, your users are experiencing strange bugs, in which users on different sides of the globe see different data. What is a likely design issue that was not accounted for when launching?

A. The system does not have Lambda Functor Repair Automations, to perform table scans and chack for corrupted partition blocks inside the Table in the recovered Region.
B. The system did not implement DynamoDB Table Defragmentation for restoring partition performance in the Region that experienced an outage, so data is served stale.
C. The system did not include repair logic and request replay buffering logic for post-failure, to re-synchronize data to the Region that was unavailable for a number of hours.
D. The system did not use DynamoDB Consistent Read requests, so the requests in different areas are not utilizing consensus across Regions at runtime.

Answer: C

Explanation:
When using multi-region DynamoDB systems, it is of paramount importance to make sure that all requests made to one Region are replicated to the other. Under normal operation, the system in question would correctly perform write replays into the other Region. If a whole Region went down, the system would be unable to perform these writes for the period of downtime. Without buffering write requests somehow, there would be no way for the system to replay dropped cross-region writes, and the requests would be serviced differently depending on the Region from which they were served after recovery. Reference:
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepI.html

NEW QUESTION 52

Your company needs to automate 3 layers of a large cloud deployment. You want to be able to track this deployment's evolution as it changes over time, and carefully control any alterations. What is a good way to automate a stack to meet these requirements?

A. Use OpsWorks Stacks with three layers to model the layering in your stack.
B. Use CloudFormation Nested Stack Templates, with three child stacks to represent the three logicallayers of your cloud.
C. Use AWS Config to declare a configuration set that AWS should roll out to your cloud.
D. Use Elastic Beanstalk Linked Applications, passing the important DNS entires between layers using the metadata interface.

Answer: B

Explanation:
Only CloudFormation allows source controlled, declarative templates as the basis for stack automation. Nested Stacks help achieve clean separation of layers while simultaneously providing a method to control all layers at once when needed.
Reference:
https://blogs.aws.amazon.com/application-management/post/TxIT9JYOOS8AB9I/Use-Nested-Stacks-to- Create-Reusable-Templates-and-Support-Role-Specialization

NEW QUESTION 55

You need the absolute highest possible network performance for a cluster computing application. You already selected homogeneous instance types supporting 10 gigabit enhanced networking, made sure that your workload was network bound, and put the instances in a placement group. What is the last optimization you can make?

A. Use 9001 MTU instead of 1500 for Jumbo Frames, to raise packet body to packet overhead ratios.
B. Segregate the instances into different peered VPCs while keeping them all in a placement group, so each one has its own Internet Gateway.
C. Bake an AMI for the instances and relaunch, so the instances are fresh in the placement group and donot have noisy neighbors.
D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughpu

Answer: A

Explanation:
For instances that are collocated inside a placement group, jumbo frames help to achieve the maximum network throughput possible, and they are recommended in this case. For more information, see Placement Groups.
Reference: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/network_mtu.htm|#jumbo_frame_instances

NEW QUESTION 57

If you're trying to configure an AWS Elastic Beanstalk worker tier for easy debugging if there are problems finishing queue jobs, what should you configure?

A. Configure Rolling Deployments.
B. Configure Enhanced Health Reporting
C. Configure Blue-Green Deployments.
D. Configure a Dead Letter Queue

Answer: D

Explanation:
Elastic Beanstalk worker environments support Amazon Simple Queue Service (SQS) dead letter queues. A dead letter queue is a queue where other (source) queues can send messages that for some reason could not be successfully processed. A primary benefit of using a dead letter queue is the ability to sideline and isolate the unsuccessfully processed messages. You can then analyze any messages sent
to the dead letter queue to try to determine why they were not successfully processed. Reference:
http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features-managing-env-tiers.html#worker-d eadletter

NEW QUESTION 62

......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:

* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently

* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff

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

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

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