

- 1) **A company hosts a web application on an Amazon EC2 instance. Users report that the web application is occasionally unresponsive. Amazon CloudWatch metrics indicate that the CPU utilization is 100% during these times. A SysOps administrator must implement a solution to monitor for this issue.**

Which solution will meet this requirement?

- A. Create a CloudWatch alarm that monitors AWS CloudTrail events for the EC2 instance.
- B. Create a CloudWatch alarm that monitors CloudWatch metrics for EC2 instance CPU utilization.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic to monitor CloudWatch metrics for EC2 instance CPU utilization.
- D. Create a recurring assessment check on the EC2 instance by using Amazon Inspector to detect deviations in CPU utilization.

- 2) **A company has an application that uses Amazon ElastiCache for Memcached to cache query responses to improve latency. However, the application's users are reporting slow response times. A SysOps administrator notices that the Amazon CloudWatch metrics for Memcached evictions are high.**

Which actions should the SysOps administrator take to fix this issue? (Select TWO.)

- A. Flush the contents of ElastiCache for Memcached.
- B. Increase the ConnectionOverhead parameter value.
- C. Increase the number of nodes in the cluster.
- D. Increase the size of the nodes in the cluster.
- E. Decrease the number of nodes in the cluster.

- 3) **A company needs to ensure that an AWS Lambda function can access resources in a VPC in the company's account. The Lambda function requires access to third-party APIs that can be accessed only over the internet.**

Which action should a SysOps administrator take to meet these requirements?

- A. Attach an Elastic IP address to the Lambda function and configure a route to the internet gateway of the VPC.
- B. Connect the Lambda function to a private subnet that has a route to the virtual private gateway of the VPC.
- C. Connect the Lambda function to a public subnet that has a route to the internet gateway of the VPC.
- D. Connect the Lambda function to a private subnet that has a route to a NAT gateway deployed in a public subnet of the VPC.

- 4) **A company runs an application on a large fleet of Amazon EC2 instances to process financial transactions. The EC2 instances share data by using an Amazon Elastic File System (Amazon EFS) file system.**

The company wants to deploy the application to a new Availability Zone and has created new subnets and a mount target in the new Availability Zone. When a SysOps administrator launches new EC2 instances in the new subnets, the EC2 instances are unable to mount the file system.

Which of the following is a possible reason for this issue?

- A. The EFS mount target has been created in a private subnet.
 - B. The IAM role that is associated with the EC2 instances does not allow the `efs:MountFileSystem` action.
 - C. The route tables have not been configured to route traffic to a VPC endpoint for Amazon EFS in the new Availability Zone.
 - D. The security group for the mount target does not allow inbound NFS connections from the security group used by the EC2 instances.
- 5) **A company uses AWS Organizations to create and manage many AWS accounts. The company wants to deploy new IAM roles in each account.**

How could a SysOps administrator deploy the new roles in each of the organization's accounts?

- A. Create a service control policy (SCP) in the organization to add the new IAM roles to each account.
- B. Deploy an AWS CloudFormation change set to the organization with a template to create the new IAM roles.
- C. Use AWS CloudFormation StackSets to deploy a template to each account to create the new IAM roles.
- D. Use AWS Config to create an organization rule to add the new IAM roles to each account.

- 6) A company runs several production workloads on Amazon EC2 instances. A SysOps administrator discovered that a production EC2 instance failed a system health check. The SysOps administrator recovered the instance manually.

The SysOps administrator wants to automate the recovery task of EC2 instances and receive notifications whenever a system health check fails. Detailed monitoring is activated for all of the company's production EC2 instances.

Which of the following is the MOST operationally efficient solution that meets these requirements?

- A. For each production EC2 instance, create an Amazon CloudWatch alarm for Status Check Failed: System. Set the alarm action to recover the EC2 instance. Configure the alarm notification to be published to an Amazon Simple Notification Service (Amazon SNS) topic.
- B. On each production EC2 instance, create a script that monitors the system health by sending a heartbeat notification every minute to a central monitoring server. If an EC2 instance fails to send a heartbeat, run a script on the monitoring server to stop and start the EC2 instance and to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. On each production EC2 instance, create a script that sends network pings to a highly available endpoint by way of a cron job. If the script detects a network response timeout, invoke a command to reboot the EC2 instance.
- D. On each production EC2 instance, configure an Amazon CloudWatch agent to collect and send logs to a log group in Amazon CloudWatch Logs. Create a CloudWatch alarm that is based on a metric filter that tracks errors. Configure the alarm to invoke an AWS Lambda function to reboot the EC2 instance and send a notification email.

- 7) The company uses AWS Organizations to manage its accounts. For the production account, a SysOps administrator must ensure that all data is backed up daily for all current and future Amazon EC2 instances and Amazon Elastic File System (Amazon EFS) file systems. Backups must be retained for 30 days.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Create a backup plan in AWS Backup. Assign resources by resource ID, selecting all existing EC2 and EFS resources that are running in the account. Edit the backup plan daily to include any new resources. Schedule the backup plan to run every day with a lifecycle policy to expire backups after 30 days.
- B. Create a backup plan in AWS Backup. Assign resources by tags. Ensure that all existing EC2 and EFS resources are tagged correctly. Apply a service control policy (SCP) for the production account OU that prevents instance and file system creation unless the correct tags are applied. Schedule the backup plan to run every day with a lifecycle policy to expire backups after 30 days.
- C. Create a lifecycle policy in Amazon Data Lifecycle Manager (Amazon DLM). Assign all resources by resource ID, selecting all existing EC2 and EFS resources that are running in the account. Edit the lifecycle policy daily to include any new resources. Schedule the lifecycle policy to create snapshots every day with a retention period of 30 days.
- D. Create a lifecycle policy in Amazon Data Lifecycle Manager (Amazon DLM). Assign all resources by tags. Ensure that all existing EC2 and EFS resources are tagged correctly. Apply a service control policy (SCP) that prevents resource creation unless the correct tags are applied. Schedule the lifecycle policy to create snapshots every day with a retention period of 30 days.

- 8) **A company is using AWS CloudTrail and wants to ensure that SysOps administrators can easily verify that the log files have not been deleted or changed.**

Which action should a SysOps administrator take to meet this requirement?

- A. Grant administrators access to the AWS Key Management Service (AWS KMS) key used to encrypt the log files.
- B. Enable CloudTrail log file integrity validation when the trail is created or updated.
- C. Turn on Amazon S3 server access logging for the bucket storing the log files.
- D. Configure the S3 bucket to replicate the log files to another bucket.

- 9) **A company is running a custom database on an Amazon EC2 instance. The database stores its data on an Amazon Elastic Block Store (Amazon EBS) volume. A SysOps administrator must set up a backup strategy for the EBS volume.**

What should the SysOps administrator do to meet this requirement?

- A. Create an Amazon CloudWatch alarm for the VolumeIdleTime metric with an action to take a snapshot of the EBS volume.
- B. Create a pipeline in AWS Data Pipeline to take a snapshot of the EBS volume on a recurring schedule.
- C. Create an Amazon Data Lifecycle Manager (Amazon DLM) policy to take a snapshot of the EBS volume on a recurring schedule.
- D. Create an AWS DataSync task to take a snapshot of the EBS volume on a recurring schedule.

- 10) **A company runs a large number of Amazon EC2 instances for internal departments. The company needs to track the costs of its existing AWS resources by department.**

What should a SysOps administrator do to meet this requirement?

- A. Activate all of the AWS generated cost allocation tags for the account.
- B. Apply user-defined tags to the instances through Tag Editor. Activate these tags for cost allocation.
- C. Schedule an AWS Lambda function to run the AWS Pricing Calculator for EC2 usage on a recurring schedule.
- D. Use the AWS Trusted Advisor dashboard to export EC2 cost reports.

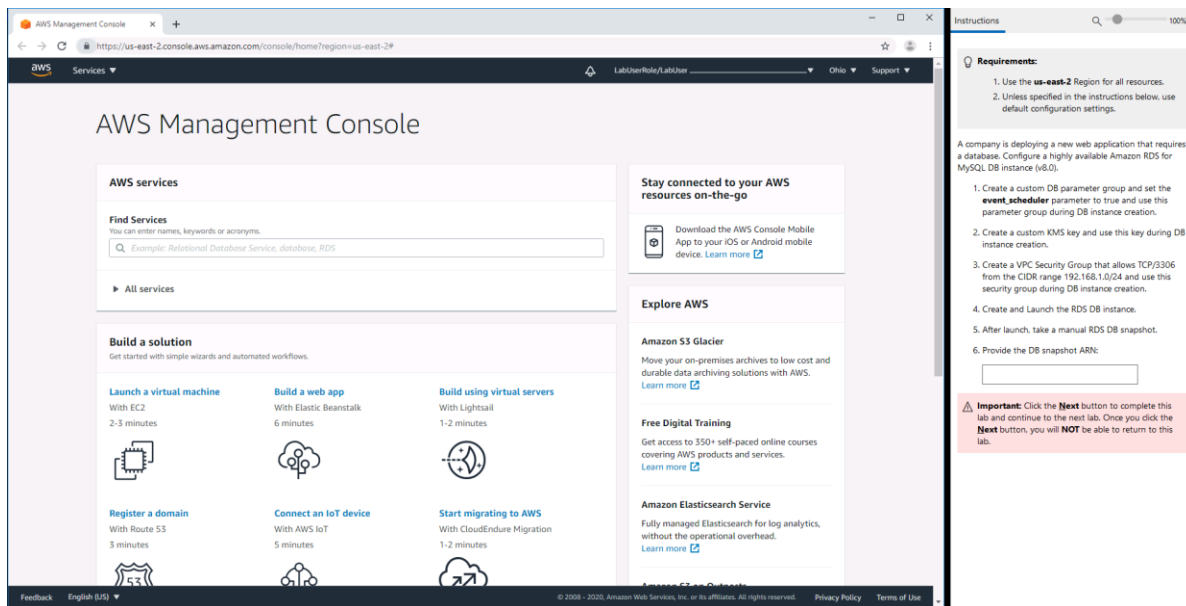
11) Sample Exam Lab

A company is deploying a new web application. Configure a highly available MySQL 8.0 database with the following:

1. Create a custom DB parameter group and set the **event_scheduler** parameter to true and use this parameter during DB instance creation.
2. Create a custom AWS Key Management Service (AWS KMS) key and use this key during DB instance creation.
3. Create a VPC security group that allows TCP port 3306 from the CIDR block 192.168.1.0/24. Use this security group during DB instance creation.
4. Launch the Amazon RDS DB instance.
5. After launch, take a manual RDS DB snapshot.

Provide the snapshot Amazon Resource Name (ARN): _____

Note: Below is a screenshot of how this sample exam lab would appear during the exam.



Answers

- 1) B — Amazon CloudWatch provides you with data and actionable insights to monitor your applications. Amazon EC2 sends metrics to CloudWatch. The CPUUtilization metric represents the percentage of allocated EC2 compute units that are currently in use on an instance. You can [create a CloudWatch alarm](#) that monitors CPUUtilization for one of your instances. For example, you might want to receive an email notification when the average CPUUtilization over a 5-minute period is greater than 75%.
- 2) C, D — The [Evictions metric](#) for Amazon ElastiCache for Memcached represents the number of non-expired items that the cache evicted to provide space for new items. If you are experiencing evictions with your cluster, it is usually a sign that you need to scale up (use a node that has a larger memory footprint) or scale out (add additional nodes to the cluster) to accommodate the additional data.
- 3) D — By default, AWS Lambda runs your functions in a secure VPC with access to AWS services and the internet. Lambda owns this VPC, which is not connected to your account's default VPC. When you [connect a Lambda function to a VPC](#) in your account to access private resources, the function cannot access the internet unless your VPC provides access. Internet access from a private subnet requires network address translation (NAT). To give your function access to the internet, route outbound traffic to a NAT gateway in a public subnet.
- 4) D — The security groups that you [associate with a mount target](#) must allow inbound access for the TCP protocol on the NFS port from the security group used by the instances.
- 5) C — With AWS CloudFormation [StackSets](#), you can create, update, or delete stacks across multiple accounts and AWS Regions with a single operation. A user in the AWS Organizations management account can create a stack set with service-managed permissions that deploys stack instances to accounts in the organization or in specific organizational units (OUs). For example, you can use AWS CloudFormation StackSets to deploy your centralized IAM roles to all accounts in your organization.
- 6) A — You can use Amazon CloudWatch alarm actions to create alarms that automatically stop, terminate, reboot, or [recover](#) your Amazon EC2 instances. For example, if an instance becomes impaired due to hardware or software issues on the physical host, loss of network connectivity, or loss of system power, you can automatically initiate a recovery action to migrate the instance to new hardware. You also can configure a message to be published to an Amazon Simple Notification Service (Amazon SNS) topic to receive a notification of the recovery action.
- 7) B — AWS Backup is a fully managed backup service that makes it easy to centralize and automate the backup of data across AWS services. The use of [tags to assign resources](#) is a simple and scalable way to back up multiple resources. Any resources with the tags that you specify are assigned to the backup plan. A [tag policy is a type of service control policy](#) (SCP) in AWS Organizations that can help you standardize and enforce tags across resources in your organization's accounts.
- 8) B — You can validate the integrity of AWS CloudTrail log files and detect whether the log files were unchanged, modified, or deleted since CloudTrail delivered them to your Amazon S3 bucket. With a validated log file, you can assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also informs

you if a log file has been deleted or changed. You gain the insight to assert positively that log files either were delivered or were not delivered to your account during a given period of time. You can [activate log file integrity validation](#) with the CloudTrail console when you create or update a trail.

- 9) C — You can use Amazon Data Lifecycle Manager (Amazon DLM) to automate the creation, retention, and deletion of Amazon Elastic Block Store (Amazon EBS) snapshots. You can [create a lifecycle policy](#) that includes specific tags to back up EBS volumes on a specified schedule and for a specified retention period. For example, you can take a snapshot of an EBS volume every day and keep the snapshots for 30 days.
- 10) B — [User-defined tags](#) are tags that you define, create, and apply to resources manually. You can use Tag Editor to search for all resources and apply tags to them. Use cost allocation tags to track your AWS costs on a detailed level. After you activate cost allocation tags, AWS uses the tags to organize your resource costs to make it easier for you to categorize and track your AWS costs. For example, to track costs by department, you can use a tag that is named "Department" with the value equal to the department name.
- 11) Lab solution:

[Create a custom DB parameter group](#) and set the `event_scheduler` parameter to true and use this parameter during DB instance creation.

- i. Open the Amazon RDS console from <https://console.amazonaws.com/rds/>.
- ii. In the **Resources** section, choose **Parameter groups**.
- iii. Choose **Create parameter group**.
- iv. In the **Parameter group family** list, select **mysql8.0**
- v. In the **Group name** box, enter the new DB cluster parameter group name of **mysql80witheventscheduler**.
- vi. In the **Description** box, enter a description for the new DB cluster parameter group.
- vii. Choose **Create**.
- viii. In the list of parameter groups, check the box next to the parameter group that you want to modify, which is **mysql80witheventscheduler**.
- ix. Choose **Parameter group actions** and choose **Edit**.
- x. In the **Filter parameters** box, enter **event_s**. This should filter just the **event_scheduler** parameter.
- xi. Choose the box for the **event_scheduler** parameter. Under **Values**, change the setting to **ON**.
- xii. Choose **Save changes**.

[Create a custom AWS Key Management Service \(AWS KMS\) key](#) and use this key during DB instance creation.

Open the AWS KMS console from <https://console.aws.amazon.com/kms>.

- i. In the navigation pane, choose **Customer managed keys**.
- ii. Choose **Create key**.
- iii. To create a symmetric CMK, for **Key type**, choose **Symmetric**.
- iv. Choose **Next**.
- v. Type the alias or display name for the CMK. For this walkthrough, use the value **mysqlDbKey** (Optional) Type a description for the CMK.
- vi. Choose **Next**.
- viii. (Optional) To add a tag, click **Add tag**. Type a tag key and an optional tag value. To add more than one tag to the CMK, choose **Add tag**.

- ix. Once completed, choose **Next**.
 - x. Select the IAM users and roles that can administer the CMK. For this walkthrough, use your IAM user.
 - xi. Choose **Next**.
 - xii. Select the IAM users and roles that can use the CMK for [cryptographic operations](#). For this walkthrough, none are needed.
 - xiii. Choose **Next**.
 - xiv. Review the key policy document that was created from your choices. Note that it can also be edited.
 - xv. Choose **Finish** to create the CMK.
- [Create a VPC security group](#) that allows TCP port 3306 from the CIDR block 192.168.1.0/24 and use this security group during DB instance creation.**
- i. Open the Amazon VPC console at <https://console.aws.amazon.com/vpc/home>.
 - ii. In the navigation pane, choose **Security Groups**.
 - iii. Choose **Create security group**.
 - iv. Enter a name for the security group (for example, **mysqlAccessGroup**) and then provide a description.
 - v. From **VPC**, select the ID of your VPC.
 - vi. Under **Inbound rules**, choose **Add rule**.
 - vii. Set **Type** to **MYSQL/Aurora**.
 - viii. Set **Source** to **My IP**.
 - ix. Scroll down and choose **Create security group**.
- [Launch the Amazon RDS DB instance](#).**
- i. Open the Amazon RDS console from <https://console.aws.amazon.com/rds/>.
 - ii. In the navigation pane, choose **Databases**.
 - iii. Choose **Create database**.
 - iv. On the **Create database** page, verify that the **Standard create** option is chosen. Then choose **MySQL**.
 - v. In the **Templates** section, choose **Production**.
 - vi. In the **DB instance identifier** section, type the name **mysqldemo**
 - vii. In the **Settings** section, set these values:
 - i. **Master password**
 - ii. **Confirm password** – Retype the password.
 - viii. In the **DB instance size** section, set these values:
 - iii. **Burstable classes (includes t classes)**
 - iv. **db.t3.micro**
 - ix. In the **Connectivity** section, for **Virtual private cloud (VPC)**, choose an existing VPC.
 - x. Expand the **Additional connectivity configuration** menu and set these values:
 - v. For **Subnet group** select the DB subnet group.
 - vi. For **Public access**, select **No**.
 - vii. For **Existing VPC security groups** choose **mysqlAccessGroup**.
 - xi. Remove the other existing security groups, such as the default security group, by choosing the **X** associated with each.
 - xii. Expand the **Additional configuration** section.
 - xiii. For the **DB parameter group**, select **mysql80witheventscheduler**
 - xiv. For **Master key**, select **mysqlDbKey**
 - xv. Choose **Create database** to create your RDS MySQL DB instance.
- After launch, [take a manual RDS DB snapshot](#).**
- i. Open the Amazon RDS console from <https://console.aws.amazon.com/rds/>.
 - ii. In the navigation pane, choose **Databases**.
 - iii. In the list of DB instances, choose the DB instance for which you want to take a snapshot.

- iv. Choose **Actions** and choose **Take snapshot**.
- v. The **Take DB snapshot** window appears.
- vi. In the **Snapshot name** box, type the name of the snapshot. For this walkthrough, use **mysqlsnapshot**.
- vii. Choose **Take snapshot**.
- viii. From the RDS console, in the navigation pane, choose **Snapshots**.
- ix. Choose the snapshot name **mysqlsnapshot**
- x. In the **Details** section, note the ARN field and the ARN.

Provide the DB snapshot ARN: _____