

# AWS Certified SysOps Administrator Practice Exam Sample Study Guide



**EVERYTHING you need from our exam experts!**

**Featuring practice questions, answers, and explanations for each question.**

**This study guide is a SAMPLE. Visit <https://aws-certifiedsysopsadministrator.examzify.com> to get the full version available exclusively to Examzify Plus pass holders .**

**Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.**

**ALL RIGHTS RESERVED.**

**No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.**

**Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.**

**SAMPLE**

## Questions

- 1. Which AWS service helps in managing user access to AWS resources?**
  - A. AWS Organizations**
  - B. AWS IAM**
  - C. AWS STS**
  - D. AWS SSO**
- 2. Why is it important to monitor S3 bucket costs?**
  - A. To ensure compliance with AWS policies**
  - B. To manage and control overall cloud spending**
  - C. To enhance data transfer speeds**
  - D. To increase data redundancy**
- 3. How can you reduce costs when using AWS services?**
  - A. By using higher-capacity instances**
  - B. By choosing AWS Free Tier options and taking advantage of reserved capacity**
  - C. By avoiding the use of Lambda functions**
  - D. By increasing your storage resources**
- 4. What solution can help reduce costs when distributing a monthly 10TB data extract?**
  - A. Keep using EFS for distribution but increase instance size**
  - B. Store the files in S3 and distribute them using a CloudFront distribution**
  - C. Use Amazon Glacier for cheaper storage**
  - D. Implement a local cache mechanism to reduce EFS load**
- 5. How can you analyze 501 errors occurring in a website hosted on S3 and served via CloudFront?**
  - A. Review the CloudFront distribution settings**
  - B. Check CloudFront access logs using Athena**
  - C. Inspect the S3 bucket policies for issues**
  - D. Reconfigure the CloudFront cache behavior**

- 6. In which AWS tool can the media company create a cost report for the S3 usage based on tags?**
- A. AWS Management Console**
  - B. AWS Cost Explorer**
  - C. AWS Resource Tagging Console**
  - D. AWS CloudWatch**
- 7. What should you implement to ensure automatic failover for an RDS database?**
- A. Create an RDS read replica in another region**
  - B. Enable Multi-AZ deployment for the RDS database**
  - C. Use AWS Lambda for real-time monitoring**
  - D. Implement manual backups daily**
- 8. How can you provide write access to users uploading their profile pictures to an S3 bucket?**
- A. Federate the users with AWS IAM**
  - B. Federate the users with Cognito**
  - C. Set S3 bucket permissions to public**
  - D. Use AWS Lambda for processing uploads**
- 9. How should permissions be configured for a user needing access to S3 buckets across multiple AWS accounts?**
- A. Create a user in the HR account**
  - B. Implement a cross-account IAM policy**
  - C. Use AWS Organizations for permissions**
  - D. Share the AWS account password**
- 10. What type of storage should the media company consider for infrequently accessed data in S3 for cost efficiency?**
- A. S3 Standard**
  - B. S3 Intelligent-Tiering**
  - C. S3 Glacier**
  - D. S3 One Zone-IA**

## **Answers**

SAMPLE

1. B
2. B
3. B
4. B
5. B
6. B
7. B
8. B
9. B
10. C

SAMPLE

## **Explanations**

SAMPLE



## 1. Which AWS service helps in managing user access to AWS resources?

- A. AWS Organizations
- B. AWS IAM**
- C. AWS STS
- D. AWS SSO

AWS Identity and Access Management (IAM) is the service specifically designed to help manage user access to AWS resources. It provides a comprehensive platform for creating and managing AWS users and groups, as well as defining permissions for those users to control what actions they can perform on AWS resources. With IAM, you can establish fine-grained access controls by creating policies that grant or deny permissions to specific AWS services and resources at the level of individual users, groups, or roles. This means you can tailor access to meet the security and operational needs of your organization, ensuring that users have the necessary permissions to perform their tasks without exposing your AWS resources to unnecessary risks. IAM also integrates well with other AWS services and is essential for managing access securely across your AWS infrastructure. It is fundamental for any organization looking to follow best practices in security and compliance while leveraging AWS services. The other options listed have useful functions in the AWS ecosystem but serve different primary purposes. For instance, AWS Organizations is primarily for managing multiple AWS accounts and consolidating billing, while AWS STS (Security Token Service) is used for temporary security credentials. AWS Single Sign-On (SSO) allows for centralized access management across multiple AWS accounts and applications but does not focus solely on the direct management of user

## 2. Why is it important to monitor S3 bucket costs?

- A. To ensure compliance with AWS policies
- B. To manage and control overall cloud spending**
- C. To enhance data transfer speeds
- D. To increase data redundancy

Monitoring S3 bucket costs is crucial for managing and controlling overall cloud spending. As organizations utilize Amazon S3 for data storage, it's easy for costs to escalate, especially with varying pricing structures based on factors like storage volume, data retrieval frequency, and data transfer operations. By regularly monitoring these costs, organizations can identify usage patterns and optimize their storage strategies to minimize expenses. For instance, they can take action to archive infrequently accessed data to lower-cost storage classes or manage the lifecycle of objects to avoid unnecessary costs. While ensuring compliance with AWS policies, enhancing data transfer speeds, and increasing data redundancy are all important considerations in a cloud setup, they do not directly address the financial aspect that monitoring costs aims to control. Cost management is critical to maintaining budgetary constraints and ensuring that the organization's cloud expenditures align with its financial strategy.

### 3. How can you reduce costs when using AWS services?

- A. By using higher-capacity instances
- B. By choosing AWS Free Tier options and taking advantage of reserved capacity**
- C. By avoiding the use of Lambda functions
- D. By increasing your storage resources

Choosing AWS Free Tier options and taking advantage of reserved capacity is a highly effective way to reduce costs when using AWS services. The AWS Free Tier allows new users to explore and experiment with a range of AWS services without incurring costs for a limited amount of usage during the first 12 months. This includes a certain number of free hours of certain compute instances, storage space, and database access, which can significantly help in reducing expenses when getting started with cloud services. Additionally, reserved capacity offers a substantial discount compared to on-demand pricing. By committing to use specific services over a one-year or three-year term, you can lock in lower rates and decrease your overall expenditure significantly. This is especially beneficial for steady-state workloads that have predictable usage patterns, as the cost savings can be quite substantial. Overall, leveraging these options not only allows users to manage budget constraints effectively but also encourages smarter usage of resources and planning in cloud computing environments.

### 4. What solution can help reduce costs when distributing a monthly 10TB data extract?

- A. Keep using EFS for distribution but increase instance size
- B. Store the files in S3 and distribute them using a CloudFront distribution**
- C. Use Amazon Glacier for cheaper storage
- D. Implement a local cache mechanism to reduce EFS load

Opting to store files in Amazon S3 and distribute them using a CloudFront distribution is an effective solution for reducing costs when distributing a large monthly data extract of 10TB. Amazon S3 is designed for high durability, availability, and cost-effective object storage. It can handle large amounts of data efficiently, and storing your data here minimizes the costs associated with traditional file systems like Elastic File System (EFS). S3 offers different classes of storage, e.g., Standard, Intelligent-Tiering, and even lower-cost options for infrequently accessed data, which can further optimize costs depending on usage patterns. Integrating CloudFront, which is Amazon's content delivery network (CDN), enhances the distribution of this data. By caching the data at edge locations, CloudFront reduces the latency for end-users while also lowering costs associated with data transfer from S3. This approach minimizes the number of requests made directly to the S3 bucket and can significantly decrease egress data transfer costs, especially when large datasets are involved. Using EFS alone, even with increased instance sizes, does not address the scalability and cost-efficiency needed for distributing 10TB of data, as EFS is generally more expensive for such large volumes and lacks the scalability of S3.

**5. How can you analyze 501 errors occurring in a website hosted on S3 and served via CloudFront?**

- A. Review the CloudFront distribution settings**
- B. Check CloudFront access logs using Athena**
- C. Inspect the S3 bucket policies for issues**
- D. Reconfigure the CloudFront cache behavior**

Analyzing 501 errors, which indicate that the server does not support the functionality required to fulfill the request, can be effectively achieved by checking CloudFront access logs using Athena. This is because the access logs provide detailed information about requests made to your CloudFront distribution, including the HTTP status codes returned to clients. Using Athena, you can query these logs to filter for the specific 501 errors. This allows you to identify patterns, such as which URL paths return 501 errors, the origins of the requests, and the associated user agents, giving you a clearer view of the context in which these errors occur. This data-driven approach helps in debugging issues effectively and making informed decisions on how to rectify them. Other methods, while potentially useful in understanding aspects of the deployment, do not directly provide insight specific to analyzing the occurrence of the error itself. The CloudFront distribution settings, for instance, would give you a broad overview of how the distribution is configured but may not directly point to the cause of 501 errors. Similarly, inspecting S3 bucket policies might reveal permission issues, but it is less likely to directly relate to 501 errors. Lastly, reconfiguring the CloudFront cache behavior could influence performance but without direct insights from the logs, it

**6. In which AWS tool can the media company create a cost report for the S3 usage based on tags?**

- A. AWS Management Console**
- B. AWS Cost Explorer**
- C. AWS Resource Tagging Console**
- D. AWS CloudWatch**

The correct choice for creating a cost report for Amazon S3 usage based on tags is found in AWS Cost Explorer. This tool is designed specifically for visualizing and analyzing cost and usage data. Cost Explorer allows users to filter and group their cost reports by various attributes, including tags assigned to S3 buckets. By utilizing tags, users can get detailed insights into costs associated with specific projects, applications, or departments, which is essential for managing and optimizing cloud spending. AWS Management Console serves as the primary user interface for accessing all AWS services, but it does not specifically focus on cost reporting in the way that Cost Explorer does. The Resource Tagging Console is primarily for managing tags and ensuring they are applied correctly to resources rather than for generating cost reports. AWS CloudWatch is used for monitoring and logging AWS resources but does not provide functionality for detailed cost analysis or reporting based on tags. Therefore, Cost Explorer stands out as the correct choice for the needs described in the question.

**7. What should you implement to ensure automatic failover for an RDS database?**

- A. Create an RDS read replica in another region**
- B. Enable Multi-AZ deployment for the RDS database**
- C. Use AWS Lambda for real-time monitoring**
- D. Implement manual backups daily**

Enabling Multi-AZ deployment for an RDS (Amazon Relational Database Service) database is the correct approach to ensure automatic failover. Multi-AZ deployment provides high availability and durability by automatically replicating the database across multiple Availability Zones. In the event of a failure in the primary zone, RDS can automatically switch to a standby instance in another zone without requiring any manual intervention. This failover process is seamless and helps maintain the availability of the database for applications that depend on it. Additionally, Multi-AZ configurations support synchronous data replication, which ensures data integrity and minimizes data loss during a failover. By fulfilling the requirement of automatic failover, Multi-AZ deployment serves critical production workloads by providing an effective disaster recovery option. Other options do not provide the same level of automatic failover capability. For example, creating a read replica in another region is mainly intended for read scalability and does not facilitate automatic failover for write operations. Using AWS Lambda for real-time monitoring, while it can help monitor the health of the database, does not provide a built-in mechanism for failover. Lastly, implementing manual backups daily is essential for data recovery but does not automatically handle failover during an outage.

**8. How can you provide write access to users uploading their profile pictures to an S3 bucket?**

- A. Federate the users with AWS IAM**
- B. Federate the users with Cognito**
- C. Set S3 bucket permissions to public**
- D. Use AWS Lambda for processing uploads**

To provide write access to users uploading their profile pictures to an S3 bucket, federating the users with Amazon Cognito is the most suitable approach. Amazon Cognito enables you to manage user registration, authentication, and access control, allowing you to create a secure environment for your users to upload content. By using Cognito, you can set up user pools and identity pools, which facilitate authentication and provide temporary AWS credentials with permissions to access certain resources, including the S3 bucket. This way, as users authenticate through Cognito, they are granted specific permissions defined in your IAM roles, allowing them to upload files to the S3 bucket without making the bucket publicly accessible. This maintains security while allowing users the required access to upload their profile pictures. Other methods listed in the options either compromise security or do not directly address the requirement of managing temporary access permissions effectively, making them less suitable for this scenario.

**9. How should permissions be configured for a user needing access to S3 buckets across multiple AWS accounts?**

- A. Create a user in the HR account**
- B. Implement a cross-account IAM policy**
- C. Use AWS Organizations for permissions**
- D. Share the AWS account password**

Implementing a cross-account IAM policy is the most effective way to configure permissions for a user needing access to S3 buckets across multiple AWS accounts. Cross-account IAM policies allow you to define permissions that explicitly grant a user or role in one AWS account access to resources in another AWS account. By using cross-account policies, you can specify the principal (the user or role from the other account), the resources they can access (such as specific S3 buckets), and the types of actions they can perform (like read, write, or delete). This method is secure and follows the least privilege principle, allowing users to have only the access necessary to perform their tasks without granting unnecessary permissions. Other options may not provide the required access or could introduce security risks. Creating a user in the HR account would limit that user's access to the resources only within that specific account and wouldn't facilitate access to other accounts' S3 buckets. Utilizing AWS Organizations for permissions would help in managing permissions across accounts within an organization but often focuses more on the structure and management of accounts rather than directly providing resource access. Sharing the AWS account password is highly discouraged; it compromises security by exposing sensitive credentials and can lead to unauthorized access.

**10. What type of storage should the media company consider for infrequently accessed data in S3 for cost efficiency?**

- A. S3 Standard**
- B. S3 Intelligent-Tiering**
- C. S3 Glacier**
- D. S3 One Zone-IA**

The most suitable storage option for infrequently accessed data, when considering cost efficiency, is S3 Glacier. S3 Glacier is specifically designed for long-term data archiving and allows users to store data at a significantly lower cost compared to other storage classes. It is ideal for data that is rarely accessed, such as backups, archives, and historical data, making it a perfect fit for a media company managing infrequently accessed assets. S3 Glacier provides a retrieval option suited for data that does not require immediate access, as retrieval times can range from minutes to hours, depending on the selected retrieval method. This trade-off between cost and retrieval time is advantageous for companies looking to minimize their storage expenses while still retaining access to archived data when necessary. Other storage classes like S3 Standard and S3 Intelligent-Tiering are better suited for frequently accessed data or data with unpredictable access patterns. S3 One Zone-IA, while also cost-effective for infrequently accessed data, lacks the level of durability and availability provided by S3 Glacier, particularly for archival purposes. Therefore, for the goal of cost efficiency in storing infrequently accessed data, S3 Glacier stands out as the optimal choice.