

Creating a Static Website for the Café

Scenario

Frank and Martha are a husband-and-wife team who own and operate a small café business that sells desserts and coffee. Their daughter, Sofia, and their other employee, Nikhil (who is a secondary school student), also work at the café. The café has a single location in a large city.

The café currently doesn't have a marketing strategy. It gains new customers mostly when someone walks by, notices the café, and decides to try it. The café has a reputation for high-quality desserts and coffees, but the café's reputation is limited to people who have visited or who have heard about it from other café customers.

Sofia suggests to Frank and Martha that they should expand community awareness of what the café has to offer. The café doesn't have a web presence yet, and it doesn't currently use any cloud computing services. However, that situation is about to change.

Lab overview and objectives

In this lab, you use Amazon Simple Storage Service (Amazon S3) to build a static website and implement architectural best practices to protect and manage your data.

After completing this lab, you should be able to do the following:

- Host a static website by using Amazon S3.
- Implement one way to protect your data with Amazon S3.
- Implement a data lifecycle strategy in Amazon S3.
- Implement a disaster recovery (DR) strategy in Amazon S3.

At the end of this lab, your architecture should look like the following example:



Note: In this challenge lab, you encounter a few tasks where step-by-step instructions are not provided. You must figure out how to complete the tasks on your own.

Duration

This lab requires approximately 60 minutes to complete.

AWS service restrictions

In this lab environment, access to AWS services and service actions might be restricted to the ones that are needed to complete the lab instructions. You might encounter errors if you attempt to access other services or perform actions beyond the ones that are described in this lab.

Accessing the AWS Management Console

1. At the top of these instructions, choose Start Lab.
 - The lab session starts.
 - A timer displays at the top of the page and shows the time remaining in the session.
Tip: To refresh the session length at any time, choose Start Lab again before the timer reaches 0:00.
 - Before you continue, wait until the circle icon to the right of the AWS link in the upper-left corner turns green. When the lab environment is ready, the AWS Details panel displays.
2. To connect to the AWS Management Console, choose the AWS link in the upper-left corner above the terminal window.
 - A new browser tab opens and connects you to the console.
Tip: If a new browser tab does not open, a banner or icon is usually at the top of your browser with a message that your browser is preventing the site from opening pop-up windows. Choose the banner or icon, and then choose Allow pop-ups.
3. Arrange the AWS Management Console tab so that it displays alongside these instructions. Ideally, you have both browser tabs open at the same time so that you can follow the lab steps.

A business request for the café: Launching a static website (challenge 1)

Sofía mentions to Nikhil that she would like the café to have a website that visually showcases the café's offerings. The website would also provide customers with business details, such as the location of the store, business hours, and telephone number.

Nikhil is happy that he has been asked to create the first website for the café.

For this first challenge, you take on the role of Nikhil and use Amazon S3 to create a basic website for the café.

Task 1: Extracting the files that you need for this lab

In this task, you extract the files that you need to create the static website.

4. Download the [.zip file](#) that you need for this lab.
5. Extract the files to your computer.

Notice that you have an index.html file and images folder that contain image files.

Task 2: Creating an S3 bucket to host your static website

In this task, you create an S3 bucket and configure it to host your static website.

6. Open the Amazon S3 console.
7. Create a bucket in the US East (N. Virginia) us-east-1 AWS Region to host your static website.
Tip: You must clear Block all public access and enable ACLs.
8. Enable static website hosting on your bucket.
Tip: You use the index.html file for your index document.

Task 3: Uploading content to your S3 bucket

In this task, you upload the static files to your S3 bucket.

9. Upload the index.html file and the CSS and images folders to your S3 bucket.
10. In a separate web browser tab, open the endpoint link for your static website.

Answering questions about the lab

Your answers are evaluated when you choose the blue Submit button at the end of the lab.

11. To access the questions in this lab, at the top of these instructions, choose AWS Details.
12. Choose the Access the multiple choice questions link.
13. On the page that loads, answer the first question:
 - Question 1: When viewing the website after Task 3, do you see the page in the browser?
14. Note: Leave the questions webpage open in your browser tab. You return to it later in this lab.

Task 4: Creating a bucket policy to grant public read access

Frank shares his plan to create many new types of pastries for the café. You realize that you need to upload an image for each new dessert that he creates and enable public access on each object. You don't want to do this process manually. Instead, you decide to create a bucket policy that automatically makes each object public when it's uploaded to the folder.

14. Create a bucket policy that grants read-only permission to public anonymous users by using the bucket policy editor.

Hint: If you get stuck, see the examples in the [AWS Documentation](#).

15. Confirm that the website for the café is now publicly accessible.

Congratulations! You now have a static website for the café.

New business requirement: Protecting website data (challenge 2)

You show Sofía the new website, and she's very impressed. Good job!

You and Sofía discuss that you will likely need to make many updates to the website as the number of café offerings expands.

Olivia, an AWS solutions architect and regular café customer, advises you to implement a strategy to prevent the accidental overwrite and deletion of website objects.

You already need to make some changes to the website, so you decide that this would be a good time to explore object versioning.

Task 5: Enabling versioning on the S3 bucket

In this task, you enable versioning on your S3 bucket and confirm that it works.

16. In the Amazon S3 console, enable versioning on your S3 bucket.

Note: Notice that after you enable versioning, you can't disable it.

17. In a text editor, open the index.html file.

For example, you could use the Notepad++ or text editor of your choice.

18. Modify the file according to the following instructions:

- Locate the first line that has the embedded CSS code `bgcolor="aquamarine"` in the HTML, and change it to `bgcolor="gainsboro"`.
- Locate the line that has the embedded CSS code `bgcolor="orange"` in the HTML, and change it to `bgcolor="cornsilk"`.
- Locate the second line that has the embedded CSS code `bgcolor="aquamarine"` in the HTML, and change it to `bgcolor="gainsboro"`.
- Save the changes.

19. Upload the updated file to your S3 bucket.

20. Reload the web browser tab with your website and notice the changes.

21. To see the latest version of the index.html file, go to your bucket and choose Show versions. You should see both versions of this file listed.

22. Return to the browser tab with the multiple choice questions for this lab, and answer the following question:

- Question 2: What is another way to ensure maximum protection and prevent the accidental deletion of a preserved version? (Hint: Review the [Amazon S3 FAQs](#).)

Architecture best practice

In this task, you used one technique to implement the architecture best practice of protecting your data.

Expand to learn more about it.

New business requirement: Optimizing costs of Amazon S3 object storage (challenge 3)

Now that you have enabled versioning, you realize that the size of the S3 bucket will continue to grow as you upload new objects and versions. To save costs, you decide to implement a strategy to retire some of those older versions.

Task 6: Setting lifecycle policies

In this task, you set a lifecycle policy to automatically move older versions of the objects in your source bucket to S3 Standard-Infrequent Access (S3 Standard-IA). The policy should also eventually expire the objects.

23. Configure two rules in the website bucket's lifecycle configuration. To receive full credit, create two separate rules. Do not configure two transitions in a single rule:
 - In one rule, move previous versions of all source bucket objects to S3 Standard-IA after 30 days.
 - In the other rule, delete previous versions of the objects after 365 days.

24. Hint: If you get stuck, see the [AWS Documentation](#) for guidance.

Note: To limit the scope of the replication to a particular bucket object (for example, the index.html file), create a tag for the object before you create the lifecycle rule.

You should now have a lifecycle configuration that moves previous versions of your source bucket objects to S3 Standard-IA after 30 days. The policy also permanently deletes the objects that are in S3 Standard-IA after 365 days.

Architecture best practice

In this task, you implemented the architecture best practice of defining data lifecycle management.

Expand to learn more about it.

New business requirement: Enhancing durability and planning for DR (challenge 4)

The next time Olivia comes to the café, you tell her about the updates to the website. You describe the measures that you took to protect the website's static files from being accidentally overwritten or deleted. Olivia tells you that cross-Region replication is another feature of Amazon S3 that you can also use to back up and archive critical data.

Task 7: Enabling cross-Region replication

In this task, you enable cross-Region replication on your source S3 bucket.

24. In a different Region than the Region for your source bucket, create a second bucket and enable versioning on it. The second bucket is your destination bucket.
25. On your source S3 bucket, enable cross-Region replication. When you create the replication rule, make sure that you do the following:
 - Replicate the entire source bucket.
 - Use CafeRole for the AWS Identity and Access Management (IAM) role. This IAM role gives Amazon S3 the permissions to read objects from the source bucket and replicate them to the destination bucket.

- If you encounter the warning The replication rule is saved, but it might not work, you can ignore it and proceed to the next step.
- On the prompt about Replicate existing objects?, choose No...

26. Hint: If you get stuck, see the [AWS Documentation](#) for guidance.

Note: CafeRole has the following permissions:

Version: 2012-10-17

Statement:

- Action:

- s3:ListBucket

- s3:ReplicateObject

- s3:ReplicateDelete

- s3:ReplicateTags

- s3:Get*

Resource:

_**

Effect: Allow

This access policy allows the role to perform the replication tasks on all S3 buckets. In a real production environment, you should restrict the policy to apply to only your source and destination S3 buckets. For more information about creating an IAM role, see [Setting Up Permissions for Replication](#).

26. Return to the browser tab with the multiple choice questions for this lab, and answer the following question:

- Question 3: Do you see the objects from your source bucket in the destination bucket?

27. Make a minor change to the index.html file, and upload the new version to your source bucket.

28. Verify that the source bucket now has three versions of the index.html file.

29. Confirm that the new object was replicated to your destination bucket. You might need to reload the browser tab.

30. Go to your source bucket, and delete the latest version of the object.

31. Return to the browser tab with the multiple choice questions for this lab, and answer the following question:

- Question 4: Was the version that you just deleted from your source bucket also deleted from your destination bucket?

Architecture best practice

In this task, you implemented the architecture best practice of automating disaster recovery.

Expand to learn more about it.

Submitting your work

32. At the top of these instructions, choose Submit to record your progress.
33. When prompted, choose Yes.
After a couple of minutes, the Grades panel appears and shows you how many points you earned for each task. If the results don't display after a couple of minutes, choose Grades at the top of these instructions.
Tip: You can submit your work multiple times. After you change your work, choose Submit again. Your last submission is recorded for this lab.
34. To find detailed feedback on your work, choose Submission Report.

Lab complete

Congratulations! You have completed the lab.

35. At the top of this page, choose End Lab, and then choose Yes to confirm that you want to end the lab.

The following message briefly appears: Ended AWS Lab Successfully