



WPI



7FACTOR

Webhooks-as-a-Service

Major Qualifying Project 2023-2024

Authored by:

Jacob Adams
Griffin Atchue
Michael Emerson
Nathan Pollock
Ryan Rabbitt

Presented to:

Professor Joshua Cuneo

Sponsored by:

Jeremy Duvall, 7Factor Founder

This report represents the work of one or more WPI undergraduate students submitted to the faculty as evidence of completion of a degree requirement. WPI routinely publishes these reports on the web without editorial or peer review.

Abstract

This project advances a Webhooks-as-a-Service (WaaS) platform, enhancing automation across software services. Key developments include Terraform deployment templates, GitHub Actions for continuous integration, and AWS S3 for hosting. The user interface has been simplified, and an embedded testing workflow enables payload verification and consumer-end receipt validation. The integration of these improvements significantly elevates the platform's scalability and usability, providing a robust and user-friendly solution for developers seeking efficient inter-service communication automation.

Table of Contents

- Table of Contents..... 3**
- List of Figures.....4**
- Introduction.....5**
 - Our Sponsor..... 5
 - Purpose..... 5
- Background..... 6**
 - Webhooks..... 6
 - Terraform..... 7
 - GitHub Actions.....8
 - Amazon Web Services (AWS).....9
 - AWS Lambda..... 9
 - AWS S3.....9
 - AWS Identity and Access Management (IAM)..... 9
- Methodology..... 11**
 - Previous Work..... 11
 - AWS Backend..... 11
 - React Frontend..... 12
 - Full Stack Result..... 13
 - Actions..... 13
 - Configurations..... 13
 - Data Transformations..... 14
 - Adopted Methodologies..... 14
 - Workflow..... 15
 - A Term..... 15
 - B Term..... 17
 - C Term..... 17
- Results..... 18**
- Recommendations..... 19**
 - Refactor Codebase..... 19
 - Customizable Transformations..... 19
 - Trigger Webhook on API Change..... 19
 - Community Hub..... 20
- Conclusion..... 20**
- References..... 21**

List of Figures

[Figure 1. Example data flow](#)

[Figure 2. Data transformation from Producer to Consumer](#)

[Figure 3. WaaS login page](#)

[Figure 4. Configuration builder UI](#)

Introduction

Our Sponsor

Our sponsor, 7Factor, is a software contracting company headquartered in Atlanta, Georgia, who creates secure and scalable software solutions by adapting to their client's needs. 7Factor makes use of small teams to handle projects more efficiently. These teams work with clients in a way that works best for them, whether that is as part of a client's in house development team, or in more independent roles. Either way, 7Factor's experienced developers work closely with clients to develop simple and effective custom built solutions maintained by 7Factor (7Factor Software, 2023).

Purpose

As with previous iterations of this project, the main goal is the continued development of an application based on a concept proposed by 7Factor. This concept involves the creation of a webhooks management tool which would allow 7Factor to easily automate workflows. Building on last year's work, this project has sought to further enhance the capabilities of a WaaS platform, providing a solution that not only allows for the fundamental one-to-many producer-consumer relationship but also simplifies the process of setting up and managing such configurations. To this end, the platform now features the ability to transform and tailor webhook payloads according to the individual schema requirements of each consumer. This ensures that as events are dispatched from a single producer, they can be seamlessly integrated into the workflows of multiple consumers without the need for manual intervention or complex coding.

Background

The implementation of our WaaS platform focuses on using Terraform deployment templates for infrastructure management, which enhances the platform's scalability by allowing for repeatable and consistent deployments. Additionally, the integration of GitHub Actions for continuous integration ensures that updates to the platform are automatically tested and deployed, maintaining a high level of code quality and platform reliability. Moreover, leveraging AWS S3 for hosting not only provides a secure and scalable cloud storage solution but also contributes to the robustness of the platform's architecture.

Significant improvements have been made to the user interface, streamlining the users' interaction with the platform and enabling more intuitive setup and configuration of webhook events. Embedded testing workflows have also been introduced, allowing users to verify payloads and validate receipt by the consumer-end, further enhancing the user experience and ensuring the reliability of inter-service communication.

Webhooks

Webhooks are a fundamental concept in modern web development, offering a powerful means for applications to communicate with each other in real-time. Putting it simply, “Webhooks let you subscribe to events happening in a software system and automatically receive a delivery of data to your server whenever those events occur” (GitHub). Conceptually, a webhook is simply an HTTP callback: when an event occurs in a source application, called a producer, a notification is sent as an HTTP POST request to a specified URL, which is processed by the receiving application, called a consumer. This mechanism facilitates an event-driven architecture, allowing producers to notify consumers of specific events without the need for polling, thereby optimizing network usage and reducing latency.

Despite their utility, managing webhooks in a scalable and efficient manner presents a challenge, particularly when moving beyond the traditional one producer to one consumer paradigm. As software systems grow in complexity, there emerges a need to dispatch events from a single producer to multiple consumers, potentially each with unique payload requirements. This is where current webhook infrastructures often fall short, as they are typically designed for direct, one-to-one communication.

To address these shortcomings, the concept of a Webhooks-as-a-Service (WaaS) platform emerges, aiming to provide a centralized solution that enables webhook management and automation at scale. WaaS is a combination of webhooks and the concept of Software-as-a-Service (SaaS). SaaS is a concept most common in cloud computing services where the provider operates, manages, and maintains the software which is accessed by paying users remotely through an online connection such as a web browser rather than by using a local program (IBM). This creates convenience for users as it enables access from any device at any time with no need to wait for updates since this is all done by the remote host. By pairing webhooks with this model, users only need to know the two endpoints of their data and can easily access their webhooks from anywhere. The necessity for such a platform is driven by the limitations of existing services in adapting to diverse and evolving consumer needs, especially when considering the transformation of payloads to fit different consumer schemas. These challenges are compounded by the more granular control over webhook workflows demanded by developers, which existing solutions may not support.

Terraform

Terraform is an open-source infrastructure as code (IaC) tool created by HashiCorp that enables developers to define and provision a datacenter infrastructure using a high-level configuration language. Infrastructure as code is a configuration management concept wherein configuration specifications are specified in documented code files so that an environment can be easily replicated and undocumented changes do not occur (Red Hat, 2022). It provides a declarative approach to infrastructure management, allowing users to specify what resources are required rather than how to create those resources. This allows for a descriptive model of infrastructure in code, which can be versioned and reused.

Terraform uses its own domain-specific language, Terraform HCL (HashiCorp Configuration Language), for defining resources, and it supports a multitude of service providers, including AWS, Google Cloud, Azure, and many more. One of Terraform's key features is its ability to manage the state of the infrastructure, keeping track of the resources it creates so that it can apply incremental changes with minimal disruption.

In this year's project, Terraform played a crucial role in automating the deployment of the WaaS platform's infrastructure. We leveraged Terraform's capability to create reproducible and

consistent environments to streamline the setup process and reduce the potential for human error. By creating Terraform deployment templates, the project ensured that resources such as AWS Lambda functions, API Gateway endpoints, and DynamoDB tables could be provisioned with the correct configurations in an automated fashion. This approach not only saved time but also enhanced the scalability of the platform by enabling quick adjustments to infrastructure based on changing requirements.

GitHub Actions

GitHub Actions is a continuous integration and continuous deployment (CI/CD) platform that allows users to automate their software build, test, and deploy pipeline directly from their GitHub repository. With GitHub Actions, users can create workflows that automatically run on specific triggers, such as pushing code to a repository or creating a pull request. Workflows are defined in YAML files within the repository, allowing for version control and code review of the CI/CD process itself.

The platform provides a rich set of prebuilt actions for common tasks, and it also allows users to create their own custom actions. GitHub Actions supports a variety of programming languages and frameworks and integrates seamlessly with GitHub's suite of tools, offering a cohesive and streamlined development experience.

In the enhancement of the WaaS platform, GitHub Actions was utilized to implement a robust CI/CD pipeline that automated testing and deployment. Each push to the repository triggered a sequence of actions that included linting code, running unit tests, and deploying to the AWS environment if those steps were successful. This ensured that any changes made to the platform's codebase were reliable and did not introduce regressions, promoting a higher standard of code quality. Furthermore, by automating these steps, the development team could focus on feature development and bug fixes without the need to manually oversee each stage of the release process. GitHub Actions thus played a pivotal role in maintaining and improving the WaaS platform, ensuring that new features and improvements were integrated smoothly and efficiently.

Amazon Web Services (AWS)

This project relied heavily on several aspects of AWS to remotely host our software to function as SaaS. AWS is a leading cloud computing platform and offers hundreds of varied services. For this project, we used several of AWS' tools in conjunction which are detailed below.

AWS Lambda

AWS Lambda is a scalable solution to programming as part of Amazon's AWS cloud services. Specifically, Lambda allows users to create and run code without the need to provision or manage computing resources directly (Vişan, 2006). Users are able to write programs referred to as lambdas which are then hosted by Amazon. Rather than the user needing to provide their own hosting service or dedicated servers, Amazon actively scales resources allocated to a given lambda so that an unused one is given no resources thus incurring no costs, while heavily used ones for major sites or applications can automatically receive as much as they need to run. This allows for a more efficient use of computing resources as well as monetary resources. Since resources are only allocated to lambdas as needed, a user only incurs costs for the computing resources used meaning a user never pays for more than they need. This makes Lambda a doubly ideal service since it both reduces costs and removes the need for us to provide the computing resources to run our functions.

AWS S3

Amazon Simple Storage Service (Amazon S3), is a cloud object and data storage solution made to be easily scalable (Amazon, 2002). As another product of AWS, Amazon S3 is compatible with the rest of the AWS suite. S3 stores data in buckets; flat, non-hierarchical data structures with each object in the bucket able to be assigned tags for easy searching and sorting. S3 buckets allow for Lambda to easily run functions on specified data making it an essential pair for Lambda based software.

AWS Identity and Access Management (IAM)

Security in cloud computing is an important issue and AWS has its own security model to address this called AWS Identity and Access Management (IAM). IAM lets customers manage

access to AWS services and resources securely. At its lowest level, AWS allows customers to assign fine grained rules called policies to control access to various areas of their services. Generally, a customer will want to employ a least privileged permissions model which aims to provide the least amount of access required for their use case. These fine-grained policies are great for least privileged security models but the difficulty lies in selecting precise policies from the multitude available.

To address this, tools like 'iamlive' have been developed, offering a method to generate tailored IAM policies. 'iamlive' is a github repository that monitors the requests you send to a cloud provider like AWS and generates a least permissions policy list automatically as you work. In our work, 'iamlive' allowed us to spend less time working with IAM policies and helped us stay focused on development.

Methodology

Previous Work

This project has been ongoing and we are now in the third year of development. We were left with a strong base to start with from the previous team's work consisting of a serverless backend utilizing AWS lambda and a React frontend. AWS lambda is a serverless computing service that allows deployment of code without having to manage or deploy one's own servers. This is organized into lambda functions, which are written by the user and are executed in response to certain events.

AWS Backend

The backend consisted of an AWS Rest API Gateway, an AWS service that allows developers to expose backend services as APIs that can be consumed by other services, linked via proxy integrations to AWS Lambda which would handle all of the logic. The structure of the backend Lambda included 4 functions as follows:

- “DB Lambda”: This function, called the “DB Lambda,” is responsible for handling all front-end requests by interfacing with the database to serve and update the configurations and actions.
- “Authorizer Lambda”: This function is used in conjunction with Auth0 to authenticate the user's JWT, a token which can be translated into JSON format that is used to verify authentication credentials, and provide access to resources.
- “Router Lambda”: This function is the endpoint called by producer applications when a webhook event is triggered. It consumes the payload from the webhook, then based on the configuration ID provided as a URL parameter, pulls the proper schema from the database to determine the destination consumers and any data transformations that need to occur. From there, it asynchronously invokes an instance of the “Child Lambda” for each consumer with the original webhook payload, the schema containing the relevant transformations, and the destination consumer URL.
- “Child Lambda”: The “Child Lambda” is invoked asynchronously by the router lambda with the original webhook payload, the schema containing the relevant transformations. It

is responsible for performing any data transformations on the payload and then forwarding the resulting payload to the provided destination consumer.

The previous team chose DynamoDB as their database to store the configurations and actions used in the applications. They chose a NoSQL database due to the high variability of the structure of the data and specifically DynamoDB due to its ease to be integrated into AWS Lambda.

React Frontend

The previous team was able to make a strong React frontend. They chose React because of its popularity and their team's familiarity with it. This front end uses typescript because their team had strong previous exposure. They also chose ant design as a component library because it offered specific components that were suited to the project. React flow is also used to help display a graphical representation of the flow between producers and consumers in their configuration editor. Auth0 is used for easy and secure authentication.

Full Stack Result

The resulting platform allowed for the creation of producer and consumer nodes and the ability to transform data between source and destination.

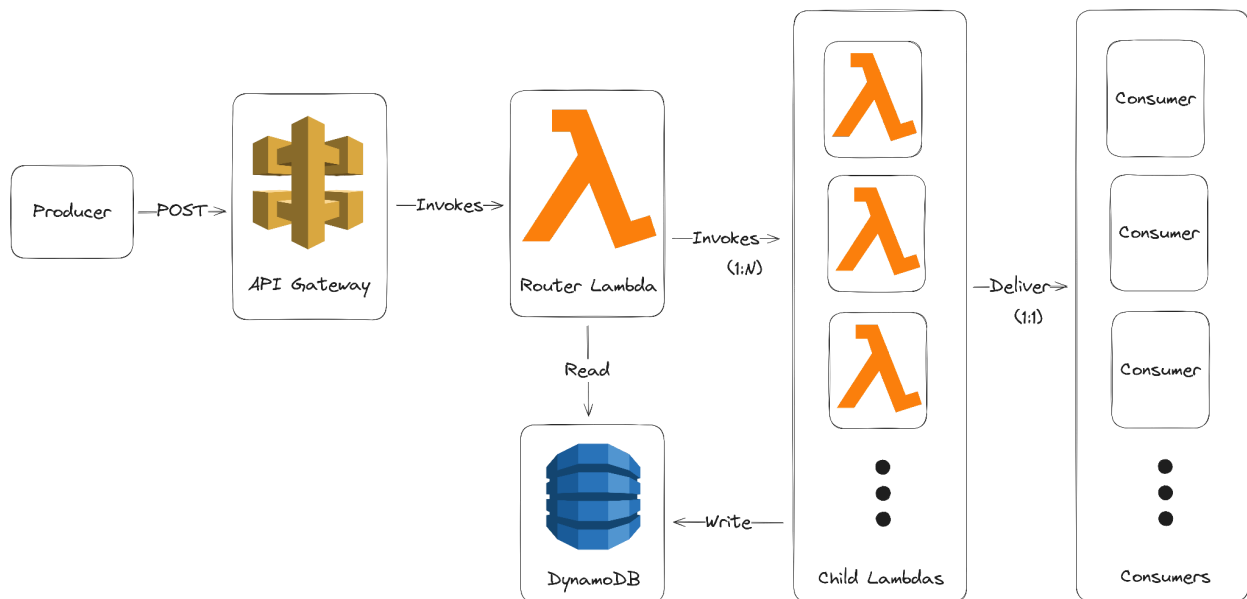


Figure 1. Example data flow

Actions

The team introduced the concept of an “action” which stores data for either a consumer or producer node. To represent a producer, the action stores fields produced by that producer. In the case of a consumer, the action will store the fields consumed by that consumer. Users are able to navigate to the “actions” tab to create, view, edit and delete actions.

Configurations

The team also introduced the concept of a “configuration” which is a mapping of a single producer action to one or more consumer actions. This configuration also stores the corresponding payload transformation details between the producer action and each consumer action. To make a full webhook configuration, the user must first have defined a producer action and at least one consumer action via the actions tab. From there, users are able to navigate to the configurations tab and create a new configuration. This opens a graphical editor where users may add their producer and consumer actions as well as define any data transformations that should occur.

Data Transformations

Data transformations define how attributes from a producer should be mapped to attributes in consumers. For example, imagine a scenario where we want to send a Discord message whenever a Github commit is made. We might want this message to include the commit message and current time. In this case, our Github producer action would have several fields including the commit message and current time, and our Discord consumer might have a single field indicating the message to send. Using data transformations, we could concatenate the commit message, current time and some nice formatting all into the message field of the Discord consumer.

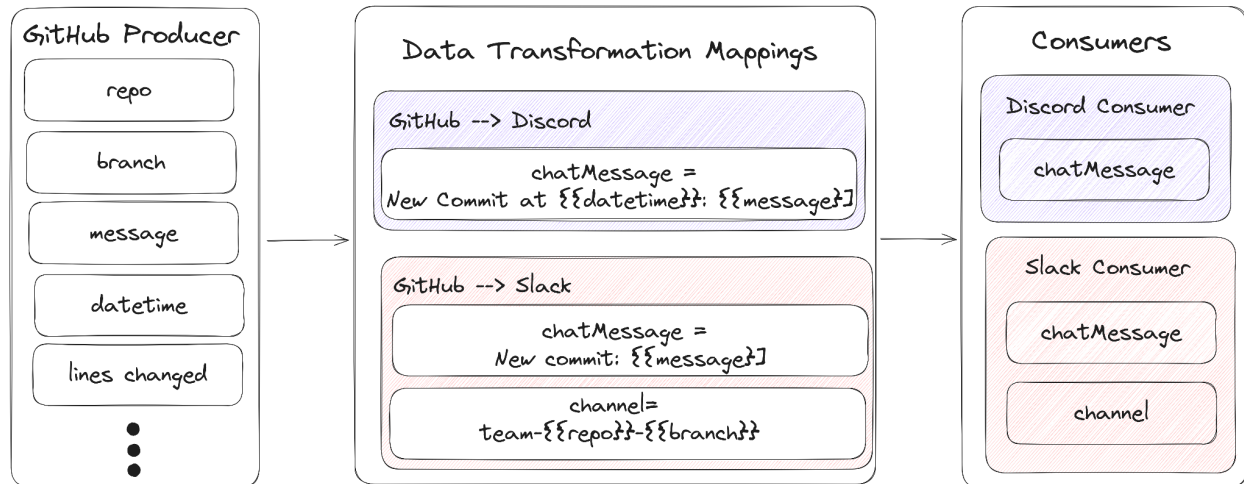


Figure 2. Data transformation from Producer to Consumer

Adopted Methodologies

The previous team's adoption of the Agile framework set a precedent for iterative development, which we continued to embrace. Their use of tools like JIRA for task management and GitHub for version control informed our decision to maintain these systems for project continuity and effective collaboration.

Workflow

Our work was split over 3 terms, and was split between our primary goals of Terraform integration, hosting the existing front end, overhauling front end design elements, and Github Actions testing workflow implementation. These goals were prioritized to our sponsor, 7Factor, during our initial pitch. Some of our other ideas included adding support for more complex data transformations between the producer and consumer, support for REST API polling, creating a developer API for external use, and a community hub for webhooks created using the 7Factor WaaS editor. These ideas were set aside in favor of our prioritized goals, as their workload would outweigh the features' usage.

A Term

After meeting with 7Factor and determining our goals for this project, the team spent a majority of A term on Terraform integration. Members of the team had previous experience with

Terraform, and were thus somewhat familiar with the tool. The development process involved iterating on different versions of the terraform template documents as the team learned how to integrate it, and finalizing the code for one section of the backend before implementing it in another section. A major challenge with our development process was finding the correct AWS permissions to execute Terraform, as there are many permissions involved to automatically deploy from Terraform.

To begin creating the terraform documents we first started by setting up the remote state management. We needed this so that we could all collaborate on the templates at once; otherwise all of our states would be different, and nothing would work. Part of creating the state management was creating a DynamoDB table as well, for the state lock.

Once we had a remote state set up we deployed the API Gateway using the swagger files, files that describe and document REST APIs, from the previous team. API Gateway allows you to import from a swagger file, so this step was pretty trivial. Once this was set up we were able to understand the structure of the API gateway and what connected to what. From here we began creating all of the API Gateway infrastructure using Terraform. This was tedious, as we needed to create API Gateway resources, methods, method responses, integrations and integration responses. A challenge we faced was getting CORS to work. CORS, or Cross-Origin Resource Sharing, is a security feature in most web browsers that is used to validate and restrict requests between origins. We had a lot of trouble with our lambda functions not returning the correct CORS headers, which prevented them from executing correctly.

Once we had all of the API Gateway created we were able to create all of the Lambda functions. These were quite simple, as we already had the code; all we had to do was create them and upload the code. Using Terraform, the lambda code was uploaded to an S3 bucket, an AWS service used as a container for static web content storage, and then moved onto the lambda.

The last piece of infrastructure we needed in the back end was the database. For this we chose DynamoDB. Again, the setup was pretty easy since we had already set it up for the state management. We just needed to create a new table and give names to some of the columns we would be using. We then had the API Gateway, database and lambdas set up, but weren't able to actually use our app without running it locally.

The previous frontend was hosted using S3. We decided to still use S3, but only to store the files for the website, and we chose to use CloudFront for the distribution of the files. In order

to set this up, we first set up the website using the console; that way we could figure out what infrastructure we would need to successfully deploy the app. Once this was done, it was reverse engineered and implemented using Terraform.

Throughout the process of deploying the infrastructure, we kept running into IAM permission issues. IAM (Identity and Access Management) permissions refer to the set of rules and policies that define what actions users, groups, and roles are allowed or denied within a system or cloud environment. Anytime we ran into a permission we didn't have, we needed to add those permissions to the IAM role, which took quite a while, considering there are so many permissions.

Once the Terraform infrastructure was in a steady state, we added GitHub actions so that any time we pushed changes to any of our repositories, it would automatically refresh the infrastructure and upload any necessary files. For the lambdas we needed to make sure that the code was reuploaded to the lambda everytime. This required using the AWS CLI, instead of just Terraform, since Terraform would see the name of the file was the same, it would have no clue that the actual code changed and needed to be re-uploaded. On the front end side of things, we needed to make sure to delete everything from the S3 bucket before running the Terraform workflow; this was for the same reason, because Terraform wouldn't know otherwise that there were changes made to the files.

Towards the end of the term, our team began the planning stage for our front end improvements. The existing Figma mockups from previous teams were used as reference and tweaked for the smaller changes, such as cleaning up the homepage/login page. New mockups were also designed for larger changes.

B Term

The group looked to move into the next stage in B Term in working on the frontend which required new code and design. To redesign the look of the frontend, our group returned to Figma, creating new mockups for the intended UI. Using the previous designs as a base, we started with the most important changes, which would help improve ease of use through increased clarification and catering to user expectations. We also made some minor changes and fixes to the front end to make the website look more professional, including adding a favicon, cleaning up the header for each page of the web app, and rewriting the "about" page to contain more general information about the project and sponsor.

We chose to redesign the UI of the configuration editor and action editor by combining them into one page, allowing for a less cluttered design. This allows the user to add a new configuration, then go to the same page to edit it. Inside the editor, the user can add and remove attributes from both the producer and consumer payloads at the same time, as well as edit the consumer's URL by clicking the "edit" option on its header. This combines all of the functionality of the two editors into one page, allowing for an easy to navigate workflow for creating and editing webhooks.

C Term

TBD

Results

With the combination of Terraform and GitHub Actions, the deployed website can be automatically updated with a successful push to the GitHub repository.

The landing page was cleaned up from its previous iterations, removing the old sample configuration builder from the authentication screen. Our team found this feature to be confusing for first time users, and decided to remove it and simplify the screen. The new authentication screen has a single box with the 7Factor logo and an authentication button, which redirects to the Auth0 login page. The top navigation bar was also redesigned, adding links to the Home and About Us pages that can be easily accessed from any screen in the web app.

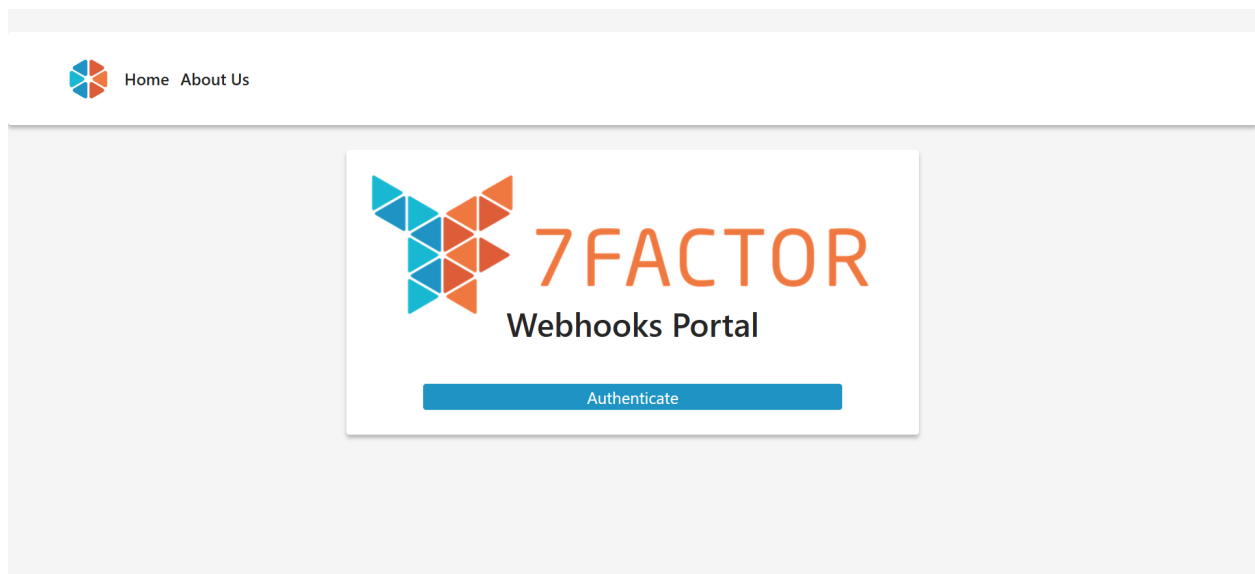


Figure 3. WaaS login page

The configuration builder was improved with multiple ease-of-use features, such as a back and save button for easier navigation, centralized editing on one page, and cleaner UI in the attributes display. Our team chose to combine the action and configuration builders into the main configuration editor, creating a more efficient webhook building editor. With the ability to edit both producers and consumers on one page, the user can avoid switching pages as much as possible. The UI of the attribute display was also improved, adding a line to show which attributes are nested in an object. This feature was taken from the previous group's Figma designs.

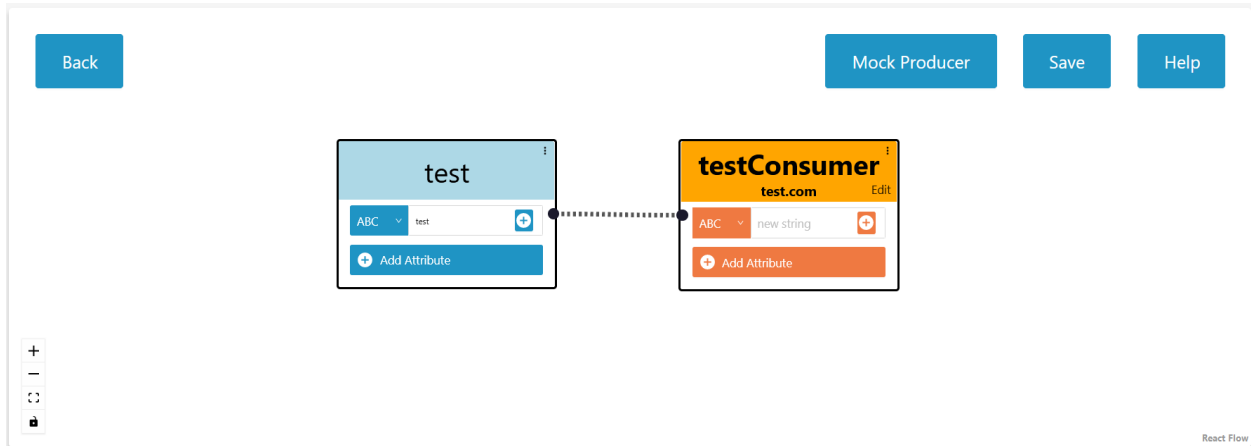


Figure 4. Configuration builder UI

Recommendations

For the purposes of furthering this project, there were some ideas we did not have the chance to implement, but would be worthwhile to implement in time.

Refactor Codebase

While the existing codebase is functioning, in its current state it poses a barrier to further development. As it is, the codebase is difficult to understand, which leads to significant amounts of time being spent deciphering code rather than making progress. Were a future group to refactor the codebase with thorough documentation this would not only ease the rest of their work, but would pave a clear path for future work.

Customizable Transformations

Adding a feature to allow users to customize data transformations between producers and consumers, such as adding a certain value to a number attribute, would allow for a broader set of use cases for this service. This feature would however require a lot of trial and error to determine which transformations are feasible for the app to handle.

Trigger Webhook on API Change

This would involve a polling system to detect a user-defined condition in a REST API. When that condition is met, a webhook is triggered.

Community Hub

As the WaaS platform is intended to be accessible for a broad audience, adding a hub where users can post their webhooks made with the service would provide a community aspect to the platform. These configurations could be downloaded and imported into the web app.

Conclusion

The work done during this project has successfully advanced the production of the WaaS platform, making several changes that will facilitate the development process for future teams. Our team implemented multiple development tools that will allow for easier deployment through AWS. The incorporation of Terraform deployment templates, GitHub actions, and AWS S3 hosting sets the stage for a streamlined development process that can be easily adopted by future teams. In addition to these implementations, a number of improvements to the UI have been made for a cleaner and more aesthetically pleasing user experience, as well as conveying the appearance of a more finished product. These improvements were made with the intention of making the web app more accessible and understandable to users. As the WaaS platform continues to evolve, we hope it will become an accessible tool that will help users more efficiently create complex webhooks for easy use.

References

- 7Factor Software. (2023, July 27). *7Factor software solutions: Customized Scalable Solutions: 7Factor*. <https://7factor.io/solutions/>.
- Amazon. (2002). *S3*. <https://aws.amazon.com/s3/>.
- Amazon. (2003). *Policies and permissions in IAM*. Amazon Web Services. https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html.
- GitHub. (n.d.). *About Webhooks*. GitHub Docs. <https://docs.github.com/en/webhooks/about-webhooks>.
- Google. "IAM Basic and Predefined Roles Reference." *Google Cloud*, cloud.google.com/iam/docs/understanding-roles.
- HashiCorp. "What Is Terraform?" *HashiCorp Developer*, developer.hashicorp.com/terraform/intro.
- IBM. (n.d.). *What is SAAS (software-as-a-service)?* <https://www.ibm.com/topics/saas>.
- Mckay, I. (2021, February 4). *IANN0036/iamlive: Generate an IAM policy from AWS, Azure, or google cloud (GCP) calls using client-side monitoring (CSM) or embedded proxy*. GitHub. <https://github.com/iann0036/iamlive>.
- Red Hat. (2022, May 11). *What is infrastructure as code (IAC)?*. Red Hat - We make open source technologies for the enterprise. <https://www.redhat.com/en/topics/automation/what-is-infrastructure-as-code-iac>.
- Vişan, S. (2006). *Lambda*. Amazon. <https://docs.aws.amazon.com/lambda/latest/dg/welcome.html>.