

Mahdi Dibaiee

Dublin, Ireland
mdibaiee@pm.me

<https://mahdi.blog>
<https://github.com/mdibaiee>
<https://linkedin.com/in/mdibaiee>

1 Experience

1.1 Estuary (Mar 2022 - Present)

Estuary Flow is a high-performance real-time DataOps platform for capturing, transforming and producing data. Worked as part of the integrations team on:

1. Develop the compatibility layer between Flow protocol and Airbyte protocol.
Impact: Allowed us to support 47 SaaS connectors from Airbyte, which drove in a large number of customers: 32% of tenants and 28% of capture tasks over one year were using these connectors.
2. Lead the evolution the Flow protocol and runtime to add an OAuth2 extension for connectors. The OAuth2 flow engages parts all the way from connectors to front-end of our application.
Impact: Ease the configuration process for our users, and ease maintenance: no need to document and troubleshoot API keys and their permission issues. The configuration-driven design allowed our support team to configure these OAuth flows without having to write code.
3. flowctl: Develop tools for easing development of, and troubleshooting of connectors:
 - Automatically & continuously inferring schema of documents coming from a connector
 - Emulating running a capture connector without the need to run an instance of Flow**Impact:** Emulation allowed for quickly iterating on a connector without having to run a connector through the whole Flow runtime, making development faster and safer.
4. Developing and maintaining real-time connectors for various technologies: Kafka, Snowflake, Big-Query, Firebolt, Cloud Firestore, Salesforce, MongoDB, Redshift, Elasticsearch, Postgres, MySQL, SQLite, PubSub, GCS, etc.
Impact: These connectors are the main driver for our customers. Most of our marketing campaigns run on these connectors and our major deals are made with customers who use these connectors.

Main Technologies: Rust, Golang, Docker, GCP

1.2 Personio (Jan 2021 - Feb 2022)

Joined as a member of the Developer Experience team, responsible for making development easier for other engineers at Personio. Personio has more than 6,000 small and medium-sized enterprise customers.

1. Surveys to understand pain points of engineers and 1:1 chats to understand needs of different teams
2. Improving deployment safety using canary deployments with Linkerd and Flagger. See: Presentation.
Impact: Over the course of three months, about 5% of deployments were rolled back automatically by our canary release process, preventing outages and protecting our SLAs. To put the figures in perspective, we had around 5 - 15 deployments per day from different teams.
3. Kubernetes Custom Resource Definitions that allowed the infrastructure teams to abstract away nuances of deployment, while exposing a basic, less error-prone interface to engineers in other teams.
4. Implementing feature flags using Split.io to allow engineers and product teams to release changes and try them out using flags for specific cohort of customers

Main Technologies: Kotlin, Kubernetes, Linkerd, Gitlab CI, AWS, Elasticsearch

1.3 AYLIEN (Jan 2017 - Dec 2022)

AYLIEN News API reads news articles from +80,000 sources, and uses machine learning models to translate and run various analysis on them, allowing for users to search across this data. There are more than 600,000 news articles processed per day.

1. Moved our infrastructure to use Infrastructure as Code solutions: Terraform, Ansible and Helm charts for Google Cloud Platform and streamline deployment process in an automatic CI/CD pipeline (previously was all manual)

Impact: Deploying resources was now driven by a CI pipeline, which made it easier for new people joining the team, made the process less human-error-prone and allowed us to more easily add or update our infrastructure resources (without worrying about permissions on GCP, etc.)

2. Incrementally moving our Ruby + Sidekiq pipeline to an actor model architecture, built with Scala + Akka.

Impact: Allowed us to capture from 65,000 new sources (5x increase) as many news sources, which translates to approximately 600,000 news articles per day. See: our blog post on this update.

3. Moved our data from our MySQL and Solr to BigTable and Elasticsearch.

Impact: Improve performance of our queries from 1.5 second 99% percentile response time to 300ms.

Main Technologies: Scala, Kubernetes, Istio, Terraform, Helm, Elasticsearch, BigTable, PubSub

2 Education

2.1 Master's: Cognitive Science

University College Dublin.

You can read my final project on <https://mahdi.blog/embodying-the-avatar-videogames/>.

2.2 Bachelor's: Computer Science

University of Science & Culture, Tehran

3 Notes

- Prior experience available on request
- Prefer remote work, but open to hybrid roles in Dublin, Ireland
- Prefer full-time roles, but open to contract roles as well
- I do not require work permit / visa sponsorship in Ireland