

Captify Engineering in a Big Nutshell

Search on its own is a highly complex data-set. Captify's technologies have been built to extract maximum value from search for brands, partners and businesses all over the world and also to innovate & improve the consumer experience.

Captify's world-class engineers, semantic specialists, and product teams are building the future of Search and as part of our Engineering Team you will play a key part in developing our offering.

The Engineering Team

The engineering team is organised into several pods, each of them is focused on a particular area or product.

Our teams are located in 3 different offices:

- 1 office in London, England
- 2 offices in Kyiv, Ukraine

How we work..

We work in an Agile environment.

As do many companies, we have adapted the methodology to our needs and the way we use it may differ from one pod to another. We may use pure Scrum (that's the case of teams in charge of applications such as C3, Sense, and Voice) or a kanban style approach for data science related projects (such as Semantics and AutoPilot).

Standups are very important for us since our engineering teams are distributed over 3 different offices, we like to catch up face to face (well, over the screen) every day.

What our Culture is like..

- We have a flexible schedule and working hours that also mean we can work from home when necessary.
- We provide a Learning and Development budget for each of our teams/members for training and conferences.
- Both our London and Kyiv offices celebrate summer and winter parties and we often try to get all of the offices together for these.
- When required by the project, our engineers may also travel to Kyiv (or travel from Kyiv to London).

Our Pod Structures

Our pods tend to be multidisciplinary: Front End, Back End, Full-stack, DevOps.

Each pod has a Team Lead whose role it is to be both the tech lead and a line manager.

Each Pod (except our QA and Infrastructure teams) have a dedicated Product Manager that provides goals, requirements, and roadmaps for our projects. The product manager is responsible for being the subject matter expert for each pod.

Our Product Pods

C3

Our C3 team is responsible for enabling the efficient management of our demand- and publisher-side activity and allowing all areas of our business access to reporting and insights.

Team structure: Fullstack and Scala developers.

Technology stack:

- React, Redux, MobX, D3, Node.js, Express, Knex, Webpack, Jest, ESLint, Prettier, Verdaccio
- Languages: Javascript, Scala

Sense

Our Sense team is responsible for providing a full-stack self-service solution to clients and partners for planning, activation, measurement, and insights, tapping into the best of Captify's proprietary data and analysis

Team structure: Fullstack and Scala developers.

Technology stack:

- React, Redux, MobX, D3, Node.js, Express, Knex, Webpack, Jest, ESLint, Prettier, Elasticsearch
- Languages: Javascript

Voice

Our voice team is building the solution to power voice search and navigation for retailers and publishers, unlocking new sources of audience insight and allowing us to better service our brand clients and publisher partners.

Team structure: Fullstack and Scala developers, ML engineers

Technology Stack:

- emotion, preact, redux-zero, Jest, enzyme, npm - yarn, lerna, node with express, redis, ESLint
- Natural Language Processing: Spacy, Flask, speech2text
- Languages: Javascript, Python

Auto Pilot

Our Auto Pilot team is responsible for fully automating the execution of campaigns and optimising towards requested KPIs and objectives.

Team structure: Data scientists and ML engineers

Technology Stack:

- Apache Spark, Spark ML, Clickhouse, Grafana, Tensorflow, HBase, PostgreSQL, Docker, Kubernetes, MLFlow, Jupyter Notebooks, AWS
- Languages: Python, Scala

Semantics

Our Semantics team are looking at deriving meaningful linguistic insight from our data in order to create a more accurate and better-performing audience. They use NLP techniques to analyse our search data.

Team structure: Data scientists and ML engineers

Technology Stack:

- Apache Spark, Spark ML, Grafana, Tensorflow, PostgreSQL, Docker, Kubernetes, Jupyter Notebooks, AWS
- Languages: Python, Scala

Big Data

Our Big Data team is enabling the optimal storage, organisation and timely retrieval of our data.

They are responsible for optimisation of ETL pipelines, maintaining over 60 Spark jobs. Building a data lake for data scientists and analysts.

Technology Stack

- Apache Spark, Apache Mesos, Akka-HTTP, Netty, PostgreSQL, Impala, Cloudera, AWS
- Languages: Scala, Python

Streaming

Our Streaming team is responsible for enabling the efficient, well-organised and robust flow of data through our systems, all the way from capture through to processing and delivery.

The team is working on real-time ingestion processes and optimisation of machine learning pipelines.

They are Processing ~100,000 events per second.

Technology stack:

- Apache Kafka, Apache Spark, Apache Flink, MLeap, Kafka-streams, Apache Mesos, Akka-HTTP, Netty, PostgreSQL, Aerospike, HBase, Elasticsearch, Clickhouse, Cloudera, AWS
- Languages: Scala, Python

Infrastructure

Our Infrastructure team are rock stars who provide and maintain all the infrastructure needed to make our technology work 24/7.

They provide high availability of services and infrastructure, monitoring clusters activity, network traffic, security, etc.

Technology stack:

- ansible, docker, kubernetes, LDAP
- Languages: bash, python, go
- Volume of data: ~1 PetaByte
- Orchestration: mesos, singularity

- Cloud: AWS
- Monitoring: Prometheus, Grafana, Graylog, Zabbix,
- Network: netflow, mikrotikOS, openvpn/ipsec

QA

The QA team are responsible for 2 main areas for automation:

1. Streaming and Big Data (QA for backend) using Scala, Specs2, Docker, Hdfs, PostgreSQL, Gatling
2. QA for Applications: JavaScript, Docker, Gatling, Cypress, Puppeteer

Who we are..

