

General

| | |
|---------------------------|---------------------------------|
| Name | Felipe Guilherme Flores e Silva |
| Day of birth | 21-03-1991 |
| Place of residence | Eindhoven, Netherlands |
| Languages | Portuguese(native) and English |
| Nationality | Brazilian |

Background

| | | |
|---------------------|-------------|--|
| Education | 2008 - 2013 | Bachelor's Degree in Computer Science - University of Londrina |
| Certificates | 2024 | Advanced Certified Scrum Developer |
| | 2019 | Professional Cloud Architect - Google Cloud Certified |
| | 2018 | Ethical Hacking and Information Security - Innovation Center Vincit |
| | 2017 | Cambridge English Level 2 Certificate in ESOL International (Advanced) |

Experiences

| | |
|----------------------------|---|
| 04/2024 - currently | Job title: Senior Software Engineer Company: Bol.com (via WAES) - Project: Fit for the Future |
| | <p>Fit for the Future(FFF) is a project to modernize Bol.com's bookkeeping system. It is built to be performant and reliable and to gracefully handle the volume and complexity of this critical area of business.</p> <p>This system consists of multiple integrations with partners (handled by individual microservices), which act as sources for a robust data pipeline (using Apache Flink) that processes, enriches, and consistently registers (BigQuery) all transactions for accounting purposes.</p> |
| Role | <p>As a contractor from WAES to Bol.com, I'm part of a Lean engineering team that includes business analysts. It means that acquiring domain knowledge and being product-driven is crucial.</p> <p>Monitoring and supporting the previous system represents 20% of the work time while developing and testing the new/cloud-native one, which is the remaining 80% of the effort. Reliability and correctness are highly praised, so the data pipeline is fully covered by e2e tests (using Cucumber).</p> |
| Highlights | <ul style="list-style-type: none">• Reliably and consistently handled millions of transactions a day.• Collaborate in a cross-functional team to deliver useful and innovative solutions• Extensively tested a high-volume data pipeline with multiple and complex flows.• Improved Code quality checkers report for all services. |
| Technologies | Google Cloud Platform (GCP), Kubernetes, Apache Flink, Google PubSub, Google BigQuery, Kotlin, Java, Grafana, Prometheus, Elastic, OpsGenie, Gitlab, Docker. |

10/2022 - 04/2024

Job title: Senior Software Engineer

Company: Wärtsilä (via WAES) - Project: Watchdogs

Watchdogs was an Incident Response Team responsible for System Reliability, monitoring, and managing production incidents.

Role

As a contractor from WAES to Wärtsilä, and leader of the Incident Response Team, I improved observability, enabling tracing, monitors, and metrics. I also conducted postmortems and followed up on actions to guarantee that lessons learned and knowledge were shared, developing runbooks to help with future troubleshooting and issue resolution.

Highlights

- Improved observability through traces, logs, monitors, and dashboards..
- Wrote run-books and troubleshooting guides.
- Lead incident postmortems and follow-up actions.
- Improved department-wide evaluation of system health and support on incidents.
- Greatly improved system observability and time to respond and resolve production incidents.

Technologies

Amazon Web Services (AWS), Kubernetes, Datadog, Apache Kafka, InfluxDB, Python, Terraform, Bamboo, Docker.

08/2021 - 10/2022

Job title: Senior Software Engineer

Company: Wärtsilä (via WAES) - Project: Expert Insights

Expert Insights (ADM) focuses on improving Marine Engine Experts' work by providing automatic anomaly detection of engine behavior, supporting preventive maintenance, saving money, and enabling safe logistics.

Expert Insights (ADM) consists of a Data Pipeline that gathers data from hundreds of sensors and feeds Machine-Learning models that learn the engine's behavior and detect possible issues.

Role

As a contractor from WAES to Wärtsilä, my responsibilities have included maintaining the pipeline, improving its performance and reliability, and implementing/translating the data scientists' research results/code into the model's lifecycle.

Highlights

- Improve system testability and reliability through continuous testing and monitoring.
- Reduced pipeline costs by redesigning critical parts, which reduced memory and computing resource usage.
- Maintained multiple high-traffic Kafka clusters for real-time processing.

Technologies

Amazon Web Services (AWS), Kubernetes, Apache Kafka, TensorFlow, Python, Java, InfluxDB, Datadog, Bamboo, Docker.

08/2020 - 08/2021

Job title: Senior Software Engineer

Company: Treez Inc. (via Truelogic) - Project: MoveTreez

MoveTreez is a system that facilitates Sales and Inventory Management of cannabis-based products from licensed distributors to licensed retailers. It keeps records and synchronizes information with the U.S. government platform Metrc (Marijuana Enforcement Tracking Reporting Compliance). Given its government/law-enforced nature, inventory tracking and management are a primary concern.

This cloud system (deployed in AWS) was developed using a microservice architecture, with some services written in Java and others in Typescript / NodeJS. PostgreSQL and Apache Pulsar backed it as a data source and message broker. Users would access a web app written in Typescript with React.

Role

As a contractor from TrueLogic to Treez, my responsibilities were developing new features and supporting inventory management services (Java/Spring). I also took part in tasks related to sales management and frontend/UI (Typescript, NodeJS, React). In this short period, there were some achievements.

Highlights

- Improved system security using OAuth2+JWT for service integration.
- Design and improve system-to-system communication synchronous (REST JSON) or asynchronous (Apache Pulsar).
- Introduced unit and integration tests to legacy and new code.

Technologies

Amazon Web Services (AWS), Java, TypeScript, React.js, Apache Pulsar, PostgreSQL, Github Actions, Circle CI, Docker.

07/2019 - 08/2020

Job title: Senior Software Engineer

Company: B3 (via Atos) - Project: Floor Plan

Floor Plan was a system that validated and tracked vehicles used as guarantees/collateral by car dealerships in bank loans. It used a pool of data both from its own and from BIN, a state-related Brazilian vehicle data center. The platform acted as a safeguard for banks and dealerships, managing contracts and enforcing legal practices between the involved parties.

The system consisted of a web server written in Java with Spring, backed by a JBoss Application Server (which we worked to migrate to Spring Boot standalone) and an Oracle database. The server had multiple integrations with internal and external services, all through APIs maintained by the team.

Role

As a contractor from Atos to B3, my responsibilities were to develop integrations with external systems while keeping API documentation (Swagger) up to date. I also advocated for developing a test culture and monitoring code health metrics.

Highlights

- Improved system security by rewriting the auth model with the OAuth2 standard.
- Advocated for and drastically improved system testability, increasing test coverage by 50 percentage points.
- Improved system stability and reliability, reducing incidents by 80%.

Technologies

Java, OracleDB, Bamboo.

| | |
|--------------------------|--|
| 06/2017 - 07/2019 | <p>Job title: Technical Lead Company: Trimble Transportation Latam - Project: Vtrips</p> <p>Vtrips was a system that provided fleet managers, drivers, and passengers with helpful information on vehicle routes, schedules, and time estimations (arrival, departure, etc.). Data were collected from vehicles via proprietary devices and from passengers and drivers via mobile phones. The system processed and showed this data in web and mobile apps as analytical data and real-time notifications.</p> <p>It was a cloud system (deployed in AWS) developed in a highly scalable microservice architecture. It was written in Java and Kotlin, using Spring and Netflix OSS. PostgreSQL, Redis, and S3 data sources and Apache Kafka for asynchronous messaging and event sourcing backed it. Users accessed a web app written in Typescript with Angular and mobile apps (maintained by other teams).</p> |
| Role | I was this project's lead developer and system architect, responsible for system design, code reviews, and standards, developing new features, and training new developers. |
| Highlights | <ul style="list-style-type: none"> • Achieved responsiveness and adaptability on the SPA web front using Angular. • Designed and implemented a scalable and efficient microservices event-driven architecture. • Enforced TDD practices, achieving 92% total system test coverage. • Reduced production incidents to virtually zero in the last five months of the project I was involved in. • The system served multiple platform consumers (web and mobile) consistently. |
| Technologies | Amazon Web Services (AWS), PostgreSQL, Apache Kafka, Redis, Kotlin, TypeScript. |
| 09/2017 - 11/2017 | <p>Job title: Technical Lead Company: Trimble Transportation Latam - Project: Velefante CLI</p> <p>Velefante was a command-line interface written in Python for DBAs to manage PostgreSQL role-based access control (RBAC). It facilitated the management of roles and users in multiple Postgres instances and databases.</p> |
| Role | I was the sole developer and maintainer. The aim was to make DBA's life easier, more secure, and less error-prone. |
| Highlights | <ul style="list-style-type: none"> • Improved security of databases, as users had only required permissions. • Improved system reliability as only specific users had read/write permissions. |
| Technologies | Python, PostgreSQL. |

| | |
|--------------------------|--|
| 01/2017 - 09/2017 | Job title: Technical Lead Company: Trimble Transportation Latam - Project: Cloud Migration |
| | <p>Cloud Migration was a project aimed at migrating every client using an on-premise server (maintained by those clients) to a cloud microservice unified solution maintained by Trimble and the cloud provider (AWS, in this case).from vehicles via proprietary devices and shown to users as reports and real-time alerts.</p> <p>It consisted of a sidecar written in Java and Spring installed along with the on-premise server. It would read its data and synchronize it to a cloud gateway system (deployed in AWS) written in Kotlin, using Spring and Netflix OSS, backed by PostgreSQL, Redis, and S3 data sources.</p> |
| Role | I was this project's lead developer and system architect, responsible for system design, code reviews, standards, and new feature development. Since it was a data-sensitive project, integrity and consistency were the main concerns, so this project had excellent test coverage from unit and integration tests on all services. |
| Highlights | <ul style="list-style-type: none">• Created an auto-updating mechanism for on-premises sidecars.• All data migration could be done safely in a couple of hours.• Data would be kept in sync without impact on the on-premise server.• Greatly improved incident response time. |
| Technologies | Amazon Web Services (AWS), PostgreSQL, RabbitMQ, Redis, Kotlin. |
| 05/2016 - 08/2016 | Job title: Full Stack Engineer Company: Trimble Transportation Latam - Project: GoGetSHP |
| | <p>GoGetSHP was a web app used internally by cartographers to extract shapefiles (.shp) of Brazilian cities from an Open Street Map database.</p> <p>The system consisted of a Go-based web server that extracted data from an OSM/PostGIS database and exported shapefiles. Users requested city shapefiles using a React web app (SPA).</p> |
| Role | I was the creator and sole developer of this tool, responsible for maintaining and improving it. My main goal was to help the cartographer team be more productive and do less repetitive work. |
| Highlights | <ul style="list-style-type: none">• Cartographer team performance increased by 10x in the following month.• Less map drawing errors were found due to the process being done a single time. |
| Technologies | Go, PostgreSQL, React.js. |

| | |
|--------------------------|--|
| 03/2015 - 01/2017 | Job title: Full Stack Engineer Company: Trimble Transportation Latam - Project: Vfleets |
| | Vfleets was a system that provided fleet managers with analytical tools and insights into telemetry and driving behavior. Data were collected from vehicles via proprietary devices and shown to users as reports and real-time alerts. |
| | It was a cloud system (deployed in AWS), developed in a highly scalable microservice architecture, written in Java, using Spring and Netflix OSS, backed by PostgreSQL, Redis, and S3 data sources, and RabbitMQ for asynchronous messaging. Users would access an AngularJS web app that got its data from REST APIs. |
| Role | I was responsible for developing new features while improving the system's capabilities, like scalability and reliability, as we would handle large amounts of telemetry data. That was a significant opportunity to learn how to deal with distributed systems problems like distributed locking, event-driven design, data partitioning, stateless services, etc. |
| Highlights | <ul style="list-style-type: none"> • Handled at least 50GB/day of raw text telemetry data. • Designed BI analytics using OLAP. • Created a table partitioning library to help manage PostgreSQL table partitions. • Secured APIs using OAuth2 standards. • Enforced test practices that significantly improved stability and reliability. |
| Technologies | Amazon Web Services (AWS), RabbitMQ, Java, PostgreSQL, Redis. |

| | |
|--------------------------|--|
| 03/2013 - 03/2015 | Job title: Junior Software Engineer Company: Trimble Transportation Latam - Project: Veltrac CS |
| | Veltrac CS was a system that provided logistics and fleet management tools, such as cost analysis, routing, telemetry, geolocation, and driving behavior. Data were collected from vehicles via proprietary devices and shown to users as reports and real-time alerts. |
| | The system consisted of an on-premises server, written in Delphi XE, with Microsoft SQLServer as a database. Users would access data via a desktop client, also written in Delphi XE. Communication between client and server was done by TCP/IP sockets using a private protocol. |
| Role | As a Junior Developer, I maintained and evolved the system. Fortunately, I also had the opportunity to learn design patterns, database design, memory management, multi-thread handling, and multiple optimization techniques. |
| Highlights | <ul style="list-style-type: none"> • Optimization of reports by rewriting SQL queries, reducing processing time and memory usage. • Advocating for Code Review and code standards reduced production incidents significantly. • Refactoring of critical parts of the system. • Improvement on the internal private framework, optimizing searches and memory management. |
| Technologies | Delphi XE, Microsoft SQLServer. |