

Leonardo Murakami

Site Reliability Engineer

contato@murakams.com [Brazil](#) github.com/leonardomurakami

Professional Summary

Passionate Site Reliability Engineer with 6+ years of experience across Data Science, MLOps, and Site Reliability Engineering. Expert in AWS cloud infrastructure, Kubernetes orchestration, and Infrastructure as Code. Strong background in building and maintaining highly available systems with focus on reliability, scalability, and performance. Currently pursuing Computer Science degree at IME-USP while working full-time.

Professional Experience



Loft

Dec 2019 - Present • 5+ years • Remote
Proptech / Real Estate Technology

Site Reliability Engineer

Feb 2024 - Present

- ✓ Implemented a unified Infrastructure as Code (IaC) repository using Terraform, introducing standardized structure, reusable modules, and best practices to accelerate infrastructure delivery. Developed internal Terraform blueprints and modules to help developers quickly provision compliant cloud resources, reducing onboarding time and configuration errors. Automated deployment workflows with Atlantis, enabling safe, collaborative infrastructure changes through pull request automation, policy checks, and streamlined approvals, significantly improving developer experience and operational reliability.
- ✓ Led the implementation of HashiCorp Vault as the primary secrets manager for all Kubernetes clusters, designing secure authentication and access policies for both applications and developers. Automated secret injection into workloads using Vault Secrets Operator and Kubernetes ServiceAccounts, reducing manual credential handling and improving compliance. Developed onboarding documentation and training sessions to enable teams to adopt secure secret management practices, resulting in faster, safer deployments and a significant reduction in security incidents related to credential exposure.
- ✓ Developed an internal engineering assistant and MCP server—an AI agent leveraging APIs from OpenAI, Claude, and others—to access internal data sources such as Jira (read-only), GitHub, and Confluence documentation. This tool streamlined information retrieval for engineers, improving productivity and decision-making by providing contextual insights and automating knowledge discovery across internal platforms.

Junior Site Reliability Engineer

May 2023 - Feb 2024

- ✓ Led the migration of over 20 microservices and internal tools from a legacy Jenkins-based CI/CD environment to a modern development Kubernetes cluster. Standardized deployment manifests and leveraged an existing GitHub Action integrated with ArgoCD to automate application delivery. Collaborated with development teams to refactor build pipelines, troubleshoot containerization issues, and ensure zero-downtime cutovers. This initiative improved infrastructure reliability, reduced deployment times, and established a scalable foundation for future growth.
- ✓ Designed and implemented a robust sandbox/internal homologation environment, empowering developers to safely validate infrastructure and application changes without impacting production. Architected automated workflows to restore production database backups into development clusters, incorporating data sanitization steps to remove or obfuscate sensitive information such as user credentials and PII. This initiative accelerated feature testing, improved developer confidence, and ensured compliance with data privacy requirements during development and QA cycles.
- ✓ Participated in weekly on-call rotations and incident response procedures.

Junior MLOps Engineer

May 2022 - May 2023

- ✓ Collaborated with data scientists and engineers to design and implement internal tools that accelerated AI development workflows. Co-developed a robust Python library that abstracted complex data retrieval, preprocessing, and model training tasks, enabling data teams to rapidly prototype and deploy machine learning models. The library featured standardized interfaces for accessing data from multiple sources, automated feature engineering utilities, and reusable model training pipelines, significantly reducing development time and improving code consistency across projects.
- ✓ Assisted in maintaining and optimizing the Amazon SageMaker environment to enable data scientists to efficiently train and test machine learning models. Implemented best practices for resource allocation and security, ensuring that users had access to scalable compute resources without bottlenecks or slowdowns. Automated environment monitoring and cost controls, provided technical support for troubleshooting training jobs, and collaborated with data scientists to streamline onboarding and improve experiment reproducibility. These efforts ensured a secure, high-performance platform for rapid model development and testing.

Junior Data Scientist

Mar 2021 - May 2022

- ✓ Created one of the company's first AI model APIs, transitioning the pricing model from a batch processing approach to real-time inference. This shift significantly increased speed and enabled new consumer experiences by allowing instant property price estimates and more dynamic product features.
- ✓ Contributed to the refactoring of existing pricing models by designing and implementing a modular framework. This new architecture separated data preprocessing, feature engineering, model training, and evaluation components, making it easier for the team to experiment with different features, algorithms, and validation strategies. Developed reusable interfaces and configuration options that allowed data scientists to quickly swap in new model types or feature sets, accelerating experimentation cycles and improving the overall robustness and maintainability of the codebase.

Data Scientist Intern

Feb 2020 - Mar 2021

- ✓ Tested the features collected during the summer internship in evaluation models to determine which ones contributed positively to predictive performance and which ones added unnecessary complexity. Conducted systematic experiments to assess the impact of each feature, helping the team focus on the most valuable data sources for real estate price modeling.
- ✓ Performed analytical studies to support the team and company in making strategic decisions. For example, conducted research to determine whether the company was pricing properties competitively in specific neighborhoods, and analyzed whether data from certain notary offices could be considered reliable when compared to others. These studies provided actionable insights that informed business strategy and data quality improvements.

Data Scientist Summer Intern

Dec 2019 - Feb 2020

- ✓ Helped collect and scrape features to support real estate pricing models, including data such as airplane routes, nearby trees, points of interest (POIs), schools, and more.

Technical Skills

Programming Languages

Python Go C/C++ Bash SQL

Cloud & Infrastructure

AWS Kubernetes Docker Terraform Helm

Monitoring & Observability

Prometheus Grafana CloudWatch Datadog
ELK Stack

Data & ML

MLOps Data Pipelines PostgreSQL BigQuery
Scikit-learn

Tools & Methodologies

Git Linux CI/CD SLI/SLO Incident Response

Education

Bachelor of Science in Computer Science

IME-USP (Instituto de Matemática e Estatística - Universidade de São Paulo) • Currently Enrolled

Pursuing degree while working full-time. Focused on algorithms, data structures, software engineering, and systems design.