

# Chris Ziemba

[cziemba@hey.com](mailto:cziemba@hey.com) • [linkedin.com/in/cziemba](https://linkedin.com/in/cziemba) • [github.com/cziemba](https://github.com/cziemba) • [chrisziemba.dev](https://chrisziemba.dev)

## Professional Work Experience

### Amazon.com – Seattle, WA

#### Senior SDE (SDE III)

*April 2016 – Present*

- Engineering team lead for the post-purchase answering CX on Rufus. Enabling Rufus to provide high quality answering for customers including after they purchase and receive (or return) a product. Guide the team to use data available at Amazon to answer questions around orders and returns, product support, and product troubleshooting.
- Enable Retrieval Augmented Generation (RAG) using millions of PDF Manuals on Amazon. This included design and build of the end-to-end pipeline to ingest, analyze, and commit the PDF contents to storage for semantic lookup during RAG on Rufus.
- Designed, deployed, and maintained an interoperability layer across all Fire TV devices, enabling seamless integration between streaming video and the Alexa platform to deliver high-performance interactive ad overlays. Implemented scalable support for features such as Add-to-Cart, Buy-Now, and Reminders via extensible API contracts, eliminating the need for frequent app updates. Developed the core authentication and authorization codepath using signature verification, enabling secure scaling to third-party applications and services.
- Engineering team lead for design and delivery of a low-touch advertising platform on multi-modal Alexa devices that enabled high profile brands to launch with rich customer experiences, examples of noteworthy brand campaigns launched include the Barbie Movie, Trolls, Coach, Kraft, and Buick.
- Re-architected and implemented a config-backed, rule-based annotation system powered by json-logic that improved efficiency from 1% to 50%, resulting in higher throughput of human-in-the-loop evaluations which result in questions being posted that are curated based on Alexa traffic patterns.
- Built an n-gram chunking and matching technique for identifying product related questions on Alexa. Leveraged a combination of in-memory Lucene indices and finite state transducers (FSTs)

that increased our product question identification accuracy by 10%, enabling us to identify 100k+ additional product questions daily.

- Serve as a Security Certifier volunteer at Amazon, collaborating with development teams to secure Amazon software through application security review processes and ensuring adherence to security-related policies. Enabled service owners to take greater ownership of their application security throughout the product development lifecycle while reducing and mitigating security risks for customers, company, partners, and employees.

## **SRC Inc. – Syracuse, NY**

### **Software Engineer**

*May 2014 – March 2016*

- Iterate on developer tools written in Python for use in encryption and messaging with CREW Duke platform.
- Support changes to configuration tools run on Windows for SRC developed hardware, as written in C#/C++.

## **Other Projects**

### **Personal Website**

*2023 – Present*

- Maintain a personal site as a scrap space to explore ideas/tech. Hosted on AWS with CDK-defined infrastructure. Static web pages are deployed to and served from S3 and server-side rendered pages are redirected to a backend hosted via my local raspberry-pi to minimize compute cost.

### **Botnek – Discord Bot**

*2020 – Present*

- Develop and maintain a discord bot that primarily provides sound effects via integration with youtube. Enables post-processing of sound effects to run on any platform that supports ffmpeg. One highlight is computation of the post-processing effect (slowdown/speedup) is never done twice using deterministic hashing for minimal latency.

### **Clarity for DotA – Android Application**

*2015 – 2016*

- At the time, the game of DotA had ~6.4 quadrillion team combinations and the app provided quick insight into suggested hero selection based on current picks. Clarity was powered by a backend hosted in AWS. An automated daily refresh of statistical data was done using match history APIs and various stats and summaries were computed/stored in a postgres database and served to the app via a public REST API endpoint. This was enabled purely by user donations and ad-revenue.

## Education

**Rochester Institute of Technology (RIT)** – Rochester, NY

*2009 – 2013*

Bachelor's in Computer Engineering

## Skills

- **Languages**
  - Most Experience: Java, TypeScript, JavaScript, Python, Bash
  - Some Experience: C++, Kotlin, Ruby, HTML/CSS, SQL, SparkSQL
  - Least Experience: C#, Go, Rust
- **AWS:** DynamoDB, S3, SQS, Elasticache (Redis), OpenSearch, ELB, Route53, ECS, SNS, CodeDeploy, StepFunctions, Lambda, Fargate, APIGateway, VPCs, RDS, CloudFormation
- **Frameworks & Libraries:** gRPC, Spring, Guice, Docker, AWS CDK, Apache Lucene, RabbitMQ, Protobuf
- **Tools:** IntelliJ, DuckDB, VIM, VSCode, Jupyter, Unix tools (git, sed, jq, awk, etc.)