

Ignateus Marten Antony Raj

Phone: 517-348-2298 | Email: ignateusmarten@gmail.com

Objective

Dynamic and passionate engineer with 12 years of experience in managing batch and real-time processing infrastructure. Experienced in creating and managing big data processing pipelines, developing models to enhance healthcare solutions, and building SaaS products in the construction technology space.

Work Experience

Co-Founder & CTO - BlueRead *(April 2026 - Present)*

Co-Founder & CTO of BlueRead, a SaaS platform to help contractors and construction managers extract data from blueprints and generate bids based on their current inventory.

- Architecting and building the full-stack platform from the ground up using Next.js on Vercel, handling blueprint ingestion, AI-powered data extraction, and automated bid generation
- Designed the authentication and multi-tenant user management system using Clerk
- Built the data layer using Prisma ORM backed by Supabase PostgreSQL for scalable relational data storage
- Implemented file upload and storage workflows using Cloudflare R2 for blueprint and document management
- Engineered an asynchronous job processing system with BullMQ, Upstash Redis, and Railway-hosted workers for background processing of blueprint analysis tasks

Software Developer - Oracle *(March 2014 - March 2026)*

Oracle Health and AI – HealthIntent Platform

- Designed and built Big Data processing pipelines handling gigabytes to terabytes of healthcare data daily across 230+ legacy data sources, integrating into a Hadoop-based ecosystem
- Engineered near real-time (NRT) data processing workflows using Apache Kafka and Apache Storm, processing millions of events per day and enabling scalable ingestion into HBase
- Recognized as the go-to engineer for real-time processing (NRT) systems, leading design decisions, troubleshooting critical issues, and ensuring successful delivery across multiple releases
- Guided and mentored engineers on real-time pipeline architecture, best practices, and debugging strategies, improving team effectiveness and reducing onboarding time
- Played a key role in developing a microservice that reduced data processing latency from ~20 seconds to ~3 seconds (~85% improvement)
- Transformed legacy data models into scalable formats optimized for high-throughput Big Data processing
- Enhanced batch and near real-time pipelines to improve data reliability, availability, and processing efficiency at scale
- Contributed to healthcare analytics and eCQM compliance initiatives, ensuring accurate and timely data for regulatory reporting
- Strengthened system reliability by improving fault tolerance and consistency across distributed Big Data systems

Cerner – Foreign System Interface (FSI)

- Designed and implemented integrations between third-party systems and native applications, enabling seamless data exchange across heterogeneous platforms
- Developed integration services primarily in Java, along with custom in-house scripting solutions to support system interoperability
- Built and maintained data transformation and interface logic to ensure accurate and consistent communication between external and internal systems
- Collaborated with cross-functional teams to onboard and integrate new external partners and systems
- Troubleshooted and resolved integration issues, improving system stability and data flow reliability

Application Programmer Co-Op - Nokia *(Spring 2013 & Fall 2013)*

- Maintained and enhanced an in-house project tracker application for an ER/SR team
- Interacted with project managers to collect requirements and implemented features to meet project needs
- Ensured the smooth operation of project tracking and contributed to improved project management

Technical Skills

Languages: Java (primary), Python, SQL, TypeScript

Distributed Systems & Big Data: Hadoop Ecosystem, Kafka, HBase, Storm, Large-Scale Data Processing

Frontend & Full-Stack: Next.js (App Router), React, Tailwind CSS, Vercel

Auth & Identity: Clerk (multi-tenant auth, session management)

Databases & ORM: Supabase PostgreSQL, Prisma ORM, PLpgSQL (stored procedures & migrations), HBase

Async Job Processing: BullMQ, Upstash Redis, Railway Workers

File Storage & CDN: Cloudflare R2 (S3-compatible object storage)

Cloud Platforms: AWS, Oracle Cloud Infrastructure (OCI), Vercel, Cloudflare, Railway

DevOps & CI/CD: CI/CD Pipelines, Spinnaker, DC/OS, Infrastructure Automation

Expertise: Real-Time & Batch Processing, Data Pipeline Design, SaaS Product Development, Scalability, Performance Optimization, Fault-Tolerant Systems

Education

MS in Computer Science

University of Texas at Dallas, Richardson, TX

Graduation: December 2013 | GPA: 3.54