Samir Thapa

Backend Engineer | Python, Node.js & AWS | Distributed Systems

Kathmandu, Nepal | +977 9861622845 | samirthapa61622845@gmail.com

linkedin.com/in/samir-thapa-73a9a31b6 | github.com/thapasamir

Summary

Backend engineer with 4 years of experience building multi-tenant SaaS platforms in Python (Django, FastAPI) and Node.js/TypeScript on AWS for US and Australian companies. I design event-driven microservices, distributed transactions and high-volume data pipelines, including a catalog sync that keeps millions of products accurate across vendors.

Technical Skills

Languages:	Python, TypeScript, JavaScript, Go, SQL
Frameworks:	Django, FastAPI, Node.js / Express, Pandas
Databases:	PostgreSQL, MySQL, MongoDB, Redis, Elasticsearch, vector databases
Cloud (AWS):	Lambda, ECS, EC2, S3, SQS, SNS, RDS, CloudWatch, CloudFront
Messaging:	AWS SQS, RabbitMQ, Celery, WebSockets, Firebase Cloud Messaging
Architecture:	Multi-tenancy, microservices, event-driven design, Saga pattern, DDD, REST APIs
AI & Search:	LLMs, embeddings, semantic and multimodal search, NLP
DevOps & Testing:	Docker, Jenkins, CI/CD, Linux, Pytest, Selenium, Behave (BDD), Locust, Django Silk

Experience

Backend Developer — A5 SYNC

Feb 2025 – Present — *United States (Remote) · Full-time*

Multi-vendor eCommerce SaaS platform where companies manage catalogs, orders and fulfillment.

Architecture & Data

- Architected the **multi-tenant backend** for a catalog of **millions of products** from multiple vendors, covering the full order lifecycle from listing and quoting to fulfillment.
- Enabled external companies to **onboard** and independently manage their own catalogs, listings, orders, quotes and fulfillment workflows, with data isolated per tenant.
- Built **data ingestion pipelines** that fetch, parse, validate and sync millions of SKUs from three vendor FTP servers, keeping product data consistent and accurate.
- Designed microservices in **Node.js (Express) and TypeScript**, using PostgreSQL for persistence, Redis for caching and RabbitMQ for asynchronous event processing.

Core Product Features

- Built the **warehouse management system** that tracks every item in stock, covering receiving, putaway, order picking and stock reservations for customers, with bulk import of inventory and locations from Excel.
- Built **subscriptions and recurring billing**: a recurring quote converts into a subscription, and the platform generates each renewal on schedule with no manual work.
- Built an **API Builder** that lets each reseller create and publish their own data APIs through configuration alone, with no code. Every API validates its results before returning them, so external consumers get consistent, reliable data.
- Built an **organization merge** feature that combines duplicate customer accounts, safely moving all quotes, orders, payments and pricing. It runs as a background job to keep the app responsive and emails account owners a summary of what changed.

Search & Integrations

- Integrated **Elasticsearch** for fast, typo-tolerant product search across millions of records, with advanced filtering and relevance tuning.
- Integrated a **payment gateway** for online transactions and built a bidirectional **HubSpot CRM** sync for sales and customer data.
- Automated product listing to **Amazon Seller Central** through its API, opening a new sales channel for vendors on the platform.

Tech: *Node.js, TypeScript, Express, PostgreSQL, Redis, RabbitMQ, Elasticsearch, HubSpot API, Amazon Seller Central API*

Backend Developer **Varicon**

Nov 2022 – May 2026 *Australia (Remote) · Part-time*

Construction-management SaaS platform covering billing, timesheets, payroll and supplier invoicing.

Architecture & Platform

- Led the migration from a **single-tenant to a multi-tenant architecture**, adding tenant-level data isolation, role-based access control (RBAC), and JWT authentication with two-factor authentication.
- Built a **supplier merge feature using the Saga pattern** to coordinate a distributed transaction across several services and databases, with transaction-state tracking and compensating rollbacks if any step failed.
- Moved billing, timesheets and third-party integrations to **event-driven microservices on AWS SQS**, decoupling them from the core application so each could scale and fail independently.
- Set up a **primary–replica database architecture with load balancing** to keep the platform responsive under high traffic.

Search & Document Processing

- Built an **OCR pipeline that learns from user corrections**: when users fix a supplier match, the correction is stored as an embedding and used to suggest the right supplier on future invoices and bills.
- Delivered a **document search and intelligence system** with semantic search over a vector database, image similarity search, LLM-generated summaries and question answering over uploaded documents.
- Created **Spotlight**, a natural-language search that lets users find product features by describing what they want to do instead of knowing the menu name.

Integrations, Infrastructure & Quality

- Built a **serverless PDF export service** on AWS Lambda that generates large reports in the background and streams progress to users over WebSockets.
- Implemented **real-time notifications** across web (WebSockets) and mobile (Firebase Cloud Messaging).
- Automated **payroll processing** with Python and Pandas, reading from Google Sheets and producing files ready to import into Xero and MYOB.
- Integrated **MYOB, Xero and EzzyBills OCR** to automate invoice and bill processing end to end.
- Set up the **testing and performance toolchain**: Selenium with Behave (BDD) for UI automation, Locust for API load testing and Django Silk for query profiling.
- Shipped production APIs with automated tests and **CI/CD pipelines** for automated deployment.

Tech: *Python, Django, PostgreSQL, Pandas, AWS (Lambda, SQS), WebSockets, vector databases, Selenium, Behave, Locust*

Education

Bachelor of Computer Application (BCA)
Hetauda City College, Tribhuvan University, Nepal