

Suresh Chandra Sekar

Senior Software Development Engineer | Backend Engineering | Python | REST APIs | Cloud Integrations
sureshkavitha1997@gmail.com | +91 7904078868 | linkedin.com/in/sureshchandrasedkar
medium.com/@sureshchandrasedkar | Chennai, India (Open to Relocation & Remote)

PROFESSIONAL SUMMARY

Senior Backend Engineer with 7+ years of Python development experience, specialising in REST API design, cloud SDK integrations, and event-driven systems. Delivered 100+ production API features across AWS, Azure, GCP, and OCI integrations using Flask and FastAPI. Strong foundation in clean code, type-safe Python, performance profiling, and backend system reliability. Comfortable working across the full backend lifecycle — from API design and integration development to debugging, optimisation, and production incident resolution.

KEY SKILLS

Languages & Frameworks: Python (Advanced), Flask (Primary), FastAPI (Intermediate), Pydantic, MSAL, Pandas, Go (Basic), Rust (Basic)

Backend Engineering: REST API Design, Multithreading & Daemons, Background Job Processing, In-memory Caching, Performance Profiling, Memory Optimisation, Python Logging

API & Messaging: Swagger / OpenAPI, gRPC, RabbitMQ, AWS SQS

Auth & Security: JWT, RBAC, Custom Role Design

Testing & Quality: pytest (basic), Clean Code, Type-Safe Python (Pydantic, type hints)

Cloud SDKs: AWS (Boto3), Azure (REST API / azure-rest-explorer), GCP (google-api-python-client, google-auth), OCI (oci)

Tool Integrations: ServiceNow CMDB, Zabbix, Docker SDK for Python

Infrastructure & IaC: Terraform (Intermediate), ARM Templates, Docker, Linux

Databases: MongoDB (primary)

DevOps & Tools: Git, Shell Scripting (Intermediate), Postman, Jenkins (deployment), Azure DevOps (ticketing)

Methodology: Agile / Scrum, Event-Driven Architecture, Microservices Architecture, Clean Code

Soft Skills: Problem Solving, Cross-team Collaboration, Technical Mentoring, Performance Optimisation

WORK EXPERIENCE

Senior Software Development Engineer | **CoreStack (formerly CloudEnablers)** Apr 2024 – Present | Chennai, India

- Designed and delivered 100+ REST API features and enhancements using FastAPI and Flask, powering cloud governance capabilities for enterprises managing 1,000+ cloud accounts across AWS, Azure, and GCP.
- Built data ingestion pipelines for the CoreStack FinOps module — extracting cost and usage data from cloud storage buckets (provided by 15+ SaaS providers including OpenAI, Anthropic, Databricks, Snowflake, Datadog, New Relic, Dynatrace, Grafana, Splunk, Elastic, Salesforce, Okta, Microsoft Teams, Dropbox, CircleCI) and transforming it into CoreStack's internal FinOps format for storage and reporting.
- Drove FinOps cost processing performance improvements and memory optimisation, reducing processing overhead and resolving production incidents for enterprise clients managing large-scale cloud and SaaS spend.
- Led budget feature revamp — redesigning the budgeting engine to support enhanced alerting, multi-dimensional budget policies, and improved accuracy of spend forecasting.
- Architected and built the StackOps Stack Pricing feature — a real-time cost estimation engine integrating third-party pricing APIs to dynamically calculate infrastructure + licence costs before provisioning, enabling data-driven deployment decisions across 1,000+ enterprise cloud accounts.
- Mentored 2 junior engineers through regular code reviews, pair programming, and guidance on Python best practices, REST API design, and cloud SDK usage (Boto3, Azure SDK).

Software Development Engineer | **CoreStack (formerly CloudEnablers)** Mar 2019 – Apr 2024 | Chennai, India

- Contributed as a pod member to a major platform revamp — responsible for migrating daemon and thread-based services to a centralised background job framework, integrating JWT authentication into assigned modules, and redesigning account onboarding flows within the project scope.
- Introduced Pydantic data models and type annotations across owned modules, improving code reliability and consistency as part of a broader codebase quality initiative across the platform.
- Built utilisation metrics support for 30+ cloud services across AWS, Azure, and OCI — architecting an event-driven discovery pipeline using AWS SQS to ingest cloud events, trigger resource discovery, and propagate updates to downstream services (ServiceNow ticket creation, utilisation fetching).
- Built CloudWatch alarm blueprints (template-based monitoring with reusable metric/condition/threshold definitions) and extended observability further by integrating Zabbix and custom AWS CloudWatch Agent (CWAgent) metrics for EC2 instances.
- Integrated AWS SSM, ARM and CloudFormation (CFN) templates, ServiceNow CMDB, and RabbitMQ — enabling secure remote resource management, automated infrastructure provisioning, ITSM ticket creation from cloud events, and asynchronous service-to-service communication via gRPC.
- Developed CoreStack-based Mistral engine templates, infrastructure automation scripts, and cloud resource provisioning workflows, building deep expertise in multi-cloud IaC from the ground up.

EDUCATION

Bachelor of Engineering — Computer Science & Engineering

2014 – 2018

T.J.S. Engineering College, Tamil Nadu (Anna University) | CGPA: 6.8

CERTIFICATIONS & CONTINUOUS LEARNING

- Claude Code in Action — Anthropic (Jan 2026)
- Master Modern Software Architecture: Microservices & Event-Driven Architecture — Udemy (Feb 2025)
- Pragmatic System Design — Udemy (Feb 2025)
- Python Unit Testing Fundamentals (unittest & pytest) — Udemy (Feb 2025)
- Clean Code — Udemy (Mar 2024)
- Python: Coding Guidelines, Tools, Tests & Packages — Udemy (Mar 2024)

NOTABLE PROJECTS

StackOps — Stack Pricing (Estimated Deployment Cost Engine) | CoreStack | Python, FastAPI, REST APIs

- Designed the API layer that calls third-party cloud pricing endpoints at deployment time — challenge was that pricing APIs across AWS, Azure, and GCP return data in different schemas, currencies, and billing units. Built a normalisation layer in Python to map all provider responses into a unified cost model (resource cost + licence/entitlement cost), then exposed a single FastAPI endpoint consumed by the frontend deployment wizard. Users see a live total cost estimate that updates as they change parameters like VM size, region, or DB tier before provisioning.

FinOps Data Ingestion Pipeline — SaaS Provider Integrations | CoreStack | Python, Flask, MongoDB, ETL

- Each of the 15+ SaaS providers (Datadog, Snowflake, Salesforce, Okta etc.) exports billing data to a cloud storage bucket in their own format — CSV, JSON, or Parquet — with different field names, date formats, and cost granularities. Built a provider-specific parser per integration that reads from the bucket, validates and cleans the data, maps it to CoreStack's internal FinOps schema, and writes to MongoDB. Designed the pipeline to be idempotent so re-runs on the same billing period do not create duplicate records.