

VISHNU PRASAD

Senior Backend and AI Engineer

Bangalore, Karnataka 560100 | +91 9880414861 | vishnuprasadapp@gmail.com | linkedin.com/in/vishapp | github.com/vishapp | vishnu-prasad.in

PROFESSIONAL SUMMARY

Senior Backend and AI Engineer with 9+ years building production-grade systems at scale. Expert in Python (FastAPI, Django, Flask), Go, AWS, Kubernetes, and Docker. Specializes in multi-agent AI systems (CrewAI, LangGraph, LangChain), retrieval-augmented generation pipelines, and LLM integration delivered at 99%+ uptime across enterprise environments. Operates as an 80% hands-on individual contributor with 20% technical leadership across distributed systems and microservices architectures.

TECHNICAL SKILLS

Languages: Python, Go, Rust, JavaScript, TypeScript

Python Frameworks & Libraries: FastAPI, Django, Flask, Celery, SQLAlchemy, Pydantic

AI and Automation: Multi-Agent Systems, RAG Pipelines, LLM Integration, Prompt Engineering, LLM Evaluation, CrewAI, LangGraph, LangChain, AutoGen, OpenAI Agents SDK

Cloud and Infrastructure: AWS, GCP, Azure, Docker, Kubernetes, Terraform, CI/CD, Jenkins, GitHub Actions

Backend Architecture: Microservices, REST APIs, GraphQL APIs, Event-Driven Architecture, Distributed Systems, TDD

Databases and Messaging: PostgreSQL, MongoDB, Redis, Kafka, RabbitMQ

Observability: Prometheus, Grafana, Distributed Tracing, Log Aggregation

Developer Tools: Claude Code, Git, Postman, Jira

NOTABLE ACHIEVEMENTS

- Shipped cross-functional backend projects serving 50+ enterprise customers at 99.9% uptime.
- Architected scalable microservices on AWS EKS achieving 40% API throughput improvement across production environments.
- Created a multi-agent AI framework that eliminated 30% of manual security investigations and cut false-positive triage time by 45%.
- Deployed a retrieval-augmented generation system processing 50,000+ security documents with 92% retrieval accuracy.
- Reduced incident response time by 50% through orchestrated playbook execution and AI-driven threat analysis.
- Scaled platforms supporting 10,000+ IoT devices with real-time data processing via Kafka and Celery.
- Cut infrastructure costs by 25% through advanced caching, query optimization, and AWS resource right-sizing.
- Produced REST and GraphQL APIs plus Python connectors for ServiceNow, Splunk, and Microsoft Sentinel, streamlining enterprise security workflows.

PROFESSIONAL EXPERIENCE

Senior Backend Developer

HeyMarvin | Bangalore, Karnataka (Remote) | Dec 2025 – Present

- Own backend services end-to-end — scoping, architecture, execution, deployment, and stakeholder communication, shipping features on tight timelines with a delivery-first mindset.
- Drive architecture and technical decisions for an AI-powered qualitative research platform, building retrieval-augmented generation pipelines for intelligent document retrieval, achieving 35% improvement in response relevance.
- Build and maintain Python (Django) microservices handling high-throughput data ingestion pipelines, including video uploads, PII in survey responses, integrations with third-party research tools, and the broader data pipeline infrastructure.
- Built a lightweight internal FastAPI service for async webhook handling and third-party research tool callbacks, using SQLAlchemy ORM for event persistence and audit logging.
- Implement integrations with third-party research and data tools, enabling seamless data flow into the platform's backend.
- Accelerated development velocity using Claude Code and Cursor IDE, streamlining code generation, reviews, and refactoring across Python codebases.
- Produced comprehensive test suites with 85%+ code coverage using pytest, ensuring reliability across microservices and API contracts.
- Identify risks early in project cycles, create mitigation plans, and communicate clearly with cross-functional stakeholders to set realistic expectations.

Senior Technical Lead

Trend Micro India Pvt Ltd | Bangalore, Karnataka | Aug 2023 – Nov 2025

- Led design and delivery of the Vision One SOAR platform backend in Python (FastAPI, Django) and Go, serving 50+ enterprise customers and reducing incident response time by 50%.
- Architected high-performance microservices on AWS EKS for the MXDR platform, achieving 40% API throughput improvement and 99.9% uptime.
- Created a multi-agent AI framework for intelligent log analysis, case prioritization, and threat hunting — eliminating 30% of manual security investigations and reducing false-positive triage by 45%.
- Constructed a retrieval-augmented generation pipeline ingesting 50,000+ security documents with vector embeddings and semantic search, achieving 92% retrieval accuracy for threat intelligence.
- Built an LLM evaluation framework benchmarking OpenAI, Anthropic, and open-source models across accuracy, latency, and cost metrics for production agent selection.
- Enforced enterprise-grade compliance modules including RBAC, audit logging, and PCI DSS standards, cutting security certification time by 60%.
- Provided billing and lifecycle automation via a subscription billing platform API and custom Python reconciliation engines, reducing finance team effort by 70%.
- Shipped REST and GraphQL APIs plus Python connectors for ServiceNow, Splunk, and Microsoft Sentinel, streamlining SIEM workflows.
- Optimized Redis caching, PostgreSQL query tuning with SQLAlchemy ORM, and AWS resource right-sizing — cut infrastructure costs 25% and improved query times 60%.
- Mentored a 6-member team on Python best practices, API design, TDD, and CI/CD standards, maintaining 95% on-time sprint delivery.

Senior Software Engineer

Trend Micro India Pvt Ltd | Bangalore, Karnataka | May 2023 – Aug 2023

- Maintained and extended backend services in Python (Django) for security operations workflows.
- Built REST APIs integrating with 5+ third-party security platforms including SIEM and EDR tools, cutting alert ingestion latency by 30% across enterprise deployments.
- Collaborated with cross-functional teams on feature delivery, bug fixes, and code reviews.
- Contributed to performance improvements and test coverage across Python microservices.

Senior Technology Lead / Technology Lead

Knowledge Lens Pvt Ltd | Bangalore, Karnataka | Jul 2018 – Apr 2023

- Launched an event-driven License Manager platform in Python (Django, Flask) for 100+ enterprise clients with Terraform and Helm deployment automation, cutting setup time by 60%.
- Engineered an IoT monitoring platform managing 10,000+ devices with real-time processing through Kafka, Celery, and Python microservices, enabling low-latency telemetry and event-driven alerting at scale.
- Developed Python-based network automation tools for protocol testing and validation, including custom packet generation, protocol state machine design, and automated diagnostics, reducing manual testing effort by 40%.
- Designed self-healing automation frameworks using Python, SSH, and Kafka to detect and recover from service failures, reducing manual support incidents by 60%.
- Planned and maintained REST APIs and backend services powering GLens and iLens analytics platforms, integrating with AWS Lambda and S3 event pipelines.
- Led Agile and Scrum practices, code reviews, and introduced unit test coverage reaching 90% across Python codebases with CI integration.

Associate Software Engineer

Tech Mahindra Pvt Ltd | Bangalore, Karnataka | Oct 2016 – May 2018

- Reduced manual data transfer effort by 50% by automating secure SFTP/FTPS exchanges via Python scripts on Linux/UNIX, eliminating error-prone manual file handling.
- Deployed system health checks and memory diagnostics using Python and shell scripting, reducing downtime by 20%.
- Managed 24x7 production incident response and root-cause analysis for enterprise clients, maintaining SLA compliance.

PROJECTS

Multi-Agent Debugger | github.com/vishapp/multiagentdebugger

Python AI package for multi-agent log analysis and API failure debugging. Supports OpenAI, Anthropic, Google, and Ollama. Production-tested across 500+ debugging sessions, cutting average debug time by 40%.

GemMate - AI Crew Orchestrator | github.com/vishapp/gemmate

Python platform for managing AI agent teams with search, file analysis, and LLM evaluation dashboard. Reduced multi-agent setup time by 60% compared to manual orchestration.

CrewAI Master | github.com/vishapp/crewaimaster

Python framework for generating and managing autonomous AI agents with CLI, memory, and RAG context retrieval. Supports rapid deployment of production-ready agent workflows.

j2toon / mask-env

Open-source Python utilities for AI tooling and developer productivity.

EDUCATION

Bachelor of Engineering (B.E.) — Mechanical Engineering

Visvesvaraya Technological University, Mysore, Karnataka | 2012 – 2016

CERTIFICATIONS

- Claude Certified Architect - Foundations — Anthropic Academy (March 2026)
- Claude with Anthropic API — Anthropic Academy (2025)
- Build a Microservice with Go — University of California, Irvine / Coursera (2022)
- Getting Started with Go — University of California, Irvine / Coursera (2022)
- Introduction to Front-End Development — Meta / Coursera (2021)
- Crash Course on Python — Google / Coursera (2020)