

Vishnu Sai

(647) 608-0741 | vishnu.sai@mail.utoronto.ca | Toronto, ON | [LinkedIn](#) | [GitHub](#) | [Portfolio](#)
Backend Engineer | AI, Infrastructure, Cloud

Education

University of Toronto, Toronto, ON Sep 2027 (Expected)
Honours Bachelor of Science, Computer Science Specialist

Experience

Cloud Administrator Co-op, CGI, Toronto, ON Feb 2026 – Present

- Authored Terraform modules provisioning Azure infra across multi-client environments, managing large-scale resources.
- Designed CI/CD pipelines (GitLab + Azure DevOps) automating deployments across dev, staging, and prod environments.
- Automated infra operations using Bash, PowerShell, Python, cutting support workload across multi-client environments.
- Monitored AKS workloads across client environments, maintaining 99.9% uptime using Azure Monitor alerting.

Software Engineering Intern, HealthcareJob.Agency, Markham, ON May 2025 – Jan 2026

- Architected staff portal from 0 to 1 serving 600+ users and 40k+ daily API requests, (Next.js, Python, PostgreSQL, Azure).
- Engineered a greedy algorithm to schedule 600+ employees, cutting assignment time by 99% (from 2 days to 15 mins).
- Designed and built real-time WebSocket task system handling 10k daily operations with RBAC and audit logging.
- Implemented RBAC authentication with HTTP-only cookies, password hashing, and rate limiting to secure the portal.
- Built a document ingestion engine supporting 35+ MIME types with malware validation and secure file processing.
- Wired hiring pipeline with applications, video interviews, and contract generation, cutting admin overhead by 70%.

Technical Project Lead, DurhamOne, Toronto, ON Oct 2025 – Present

- Led technical delivery of a CMS platform for a nonprofit, launching a production website serving 200+ users; coordinated a 3-developer team and delivered 20+ features from ticket to production.
- Introduced platform patterns (caching, sync-and-serve, report storage) removing runtime deps and improving reliability.

Software Developer, Amicare, Toronto, ON Sep 2025 – Feb 2026

- Built a FastAPI backend for a multi-tenant healthcare platform, normalizing NoSQL data and integrating Google Maps APIs; refactored synchronous matching logic to async to improve throughput under load.
- Engineered institutional filtering and admin chat interfaces with secure data segregation across care providers.

Projects

BeePrepared - 1st Place, Microsoft + Kyndryl Azure Track (Jan 2026) [Demo](#) | [Devpost](#) | [GitHub](#)

- Implemented backend for edtech app that converted lecture material into study kits, reducing prep time by 95%.
- Built an async ingest/generate pipeline using FastAPI and DAG orchestration to process mp4, pdf, and pptx lecture materials through AI workflows.
- Architected a job queue pipeline decoupling AI generation from the UI, preventing request timeouts under load.

Wisp [Devpost](#) | [GitHub](#)

- Built a local-first file manager with a Celery job pipeline that scans and indexes files for semantic search, storing embeddings in LanceDB and metadata in SQLite.
- Solved the FastAPI/Celery async boundary by isolating event loops in separate threads; implemented exponential backoff for SQLite lock contention under rapid concurrent writes.

Lumen - 1st Place, Best Use of MongoDB Atlas (Aug 2025) [Demo](#) | [Devpost](#) | [GitHub](#)

- Reduced AI audio response latency by 80% (20s → 4s) by migrating from text based LLM to Gemini Live streaming.
- Built MongoDB aggregation pipelines analyzing mood trends across 7 metrics for 1000+ records.

Skills

Languages: Python, TypeScript, JavaScript, SQL, Bash, PowerShell, Java, C

Backend: FastAPI, Node.js, Express.js, WebSockets, Celery, Redis, REST, Async/Concurrency

Infrastructure: Terraform (IaC), Docker, Kubernetes, GitLab CI/CD, Azure, AWS, Linux

Data: PostgreSQL, MongoDB, SQLite, LanceDB (Vector DB), ETL Pipelines

Certifications

Microsoft Certified: Azure Fundamentals (AZ-900) - Credential ID: E77F0DE5672AE6D - Certification number: 1005F8-361CA9 - [Verify Credential](#)