

# Md. Rakibul Hasan

+880 1819616939 • hello@pesnik.dev • pesnik.dev • in pesnik • pesnik  
ORCID HuggingFace Kaggle PyPI

## Professional Summary

---

Lead Engineer with 6+ years across software engineering, distributed data systems, and big data platforms — now specializing in enterprise GenAI architecture. Architecting LLMops pipelines, RAG knowledge retrieval systems, and autonomous AI agent frameworks for the past 2 years. Hands-on in open-source LLM deployment, QLoRA fine-tuning (LoRA/PEFT), and GPU-based WebGPU inference. Active open-source contributor (PyPI, HuggingFace, 200+ GitHub repos); technical mentor across AI excellence initiatives. Evaluated and adopted emerging GenAI frameworks ahead of enterprise adoption curves (MCP, WebGPU inference, QLoRA fine-tuning).

## Experience

---

### Lead Engineer, Enterprise Data Engineering: Banglalink Digital, Dhaka, BD

- Sole data platform engineer and DevOps owner across full stack: IPDR billions-of-rows ETL and near-real-time microbatch pipeline for unplanned volume spikes, enterprise analytics on Kubernetes, and Airflow orchestration — 6 years continuous infrastructure ownership (Kubernetes, Docker Swarm, Helm, CI/CD, n8n, CCE)
- Architected enterprise **GenAI agent orchestration** platform (IT Toolkit: Tauri 2/Rust + MCP client/server) with multi-agent tool-calling, autonomous workflow agents, context/memory management, and AI observability; profiled for Slack, Jira, ServiceNow, Okta, M365
- Co-architected enterprise **RAG pipeline** (LangChain, LlamaIndex) for internal knowledge retrieval; designed semantic search with embeddings, hallucination mitigation, vector store integration, and AI guardrails; mentored GenAI engineering team
- Built **LLMOps** pipeline for Day-2 Ops: autonomous AI observability, model evaluation hooks, LLM-powered L1 engineering assistant with prompt management and self-healing ETL
- Designed Agent Sandbox — AI agent evaluation environment with responsible AI controls, AI governance guardrails, and live browser/CLI access for agentic testing
- Architected AI-ready **distributed data platform** (Data Agent Ready Warehouse) with vector database-compatible layer for semantic search and embedding workloads
- Built Migration Hub: GenAI-powered web UI with LLM-assisted SQL translation and prompt engineering for Teradata → Huawei DWS migration at scale
- Deployed enterprise self-service analytics platform (Apache Superset fork) on Kubernetes (solo DevOps, 100+ users); built superset-hetuengine-connector (open-source, PyPI — first-of-kind Superset → Huawei HetuEngine connector)
- Solo migrated tier-1 service Airflow 2.6 → 3.1.8 with zero unplanned downtime
- Coached teams on Python and automation; self-hosted n8n; built macOS VPN client (tigerstripes-vpn); delivered yet-another-course Python training program

### Big Data Service Lead Engineer: Banglalink Digital, Dhaka, BD

- Deployed ML models (XGBoost) on CDP Cluster with MLOps practices; resolved source-level integration issues; engineered SIM lifecycle system (Spring Boot + MySQL)
- Developed PySpark distributed workflows; modernized network transmission with graph algorithms; built Rust/React desktop app for CSV validation
- Built real-time streaming pipeline with **Kafka** and Apache Ignite for high-throughput event processing on the big data platform

### Software Engineer: Trucklagbe, Dhaka, BD

- Built API Gateway (Express/Nginx, 20+ microservices); Selenium Grid CI/CD (60% QA reduction); MITM Proxy inspector; business automation (Node.js, Google Apps Script)
- Designed project versioning architecture for Node.js and Beego backends using Git submodules — resolved multi-project versioning challenges across 20+ services
- Analysed MongoDB mobile app user behaviour — categorised user actions, surfaced usage patterns, and presented findings via Metabase dashboards to inform product decisions

### Junior Software Developer: Trucklagbe, Dhaka, BD

- GPS parser integrating 1500+ active devices; SaaS REST APIs for enterprise clients; cash disbursement module (Node.js, MySQL, Angular)

### Software Engineering Consultant: Freelance / Vendor, Remote

- Delivered production platforms: c-onelogistics.com, brogrammerslab.tech, returncarbd.com, and others
- Built real-time Firebase/JS chat application for a Bangladeshi delivery service — coordinating riders, admins, and customers across live order flows
- Automated Statefarm.com web workflows using Java and Selenium — streamlined repetitive processes as a vendor engagement

## Key Projects

---

**Enterprise GenAI Agent Orchestration Platform (IT Toolkit): Tauri 2, Rust, Next.js 15, MCP, Fluent UI v9**  
[github.com/pesnik/ittoolkit](https://github.com/pesnik/ittoolkit)

Multi-agent orchestration (MCP client/server), CLI/browser/computer-use tool hierarchy, v2 autonomous workflow agents engine (auto/agent/human actors), context/memory management, AI guardrails, AI observability; enterprise site profiles (Slack, GitHub, Jira, ServiceNow, Okta)

**Personal AI Agent — QLoRA Fine-tuned LLM: Python, PEFT, TRL, BitsAndBytes, HuggingFace** · [github.com/pesnik/fine-tune](https://github.com/pesnik/fine-tune)

End-to-end **LoRA/PEFT** fine-tuning: TinyLlama 1.1B, 4-bit QLoRA (BitsAndBytes), LoRA rank-32, SFTTrainer, custom dataset curation; published as [pesnik/pesnik-lora](https://github.com/pesnik/pesnik-lora) + [pesnik/pesnik](https://github.com/pesnik/pesnik) on HuggingFace; live GPU-based inference via WebGPU at [pesnik.dev](https://pesnik.dev)

**Enterprise RAG Knowledge Retrieval System: LangChain, LlamaIndex, Python, Vector Store**

Enterprise **RAG pipeline**: semantic search, embeddings, vector database integration, hallucination mitigation, AI guardrails, and context/memory management for Banglalink internal knowledge retrieval

**LLMOps & AI Observability Platform (DAGnostics): Python, FastAPI, LLMs, Drain3** · [github.com/pesnik/diagnostics](https://github.com/pesnik/dagnostics)

Production **LLMOps** monitoring: LLM-based ETL failure analysis, Drain3 log clustering, anomaly detection, AI observability dashboards, model evaluation and alerting pipeline

**Migration Hub & tpt-builder — Teradata to Huawei DWS: TypeScript, Python, LLM, PyPI**

GenAI-powered prompt engineering platform for LLM-assisted SQL dialect translation and query validation; **tpt-builder** (published on PyPI) generates Teradata PT scripts programmatically — core open-source tooling for enterprise DWH migration

**Informatica-to-Airflow DAG Converter: TypeScript, Oracle**

VS Code extension automating ETL migration; adopted by multiple engineers, saving 3+ hours per workflow

**DAGSonar — Airflow Task Change Visibility: Python, Apache Airflow, PyPI** · [github.com/pesnik/dagsonar](https://github.com/pesnik/dagsonar)

Tracks DAG modifications, version diffs, and configuration changes across Airflow pipelines; delivers real-time notifications to engineering teams; published as [dagsonar](https://github.com/pesnik/dagsonar) on PyPI

**PND Service BD — B2B/B2C E-commerce Platform: Laravel, MySQL, PHP, Docker**

Full-stack e-commerce platform serving 6,000+ active users; payment gateway integration, inventory management, and order tracking; deployed on DigitalOcean with Docker

**cocoon — Secure Password Manager: Rust** · [github.com/pesnik/cocoon](https://github.com/pesnik/cocoon)

Memory-safe local password vault built in Rust; encrypted storage with zero external dependencies; demonstrates security-first systems design and Rust ownership model

## Technical Skills

---

**Languages:** Python, Java, TypeScript, JavaScript, Rust, Golang, SQL, C++, PHP, Shell, Swift

**GenAI / LLMOps:** LLMOps, RAG Pipelines, Agent Orchestration, Prompt Engineering, Semantic Search, Vector Databases, Embeddings, AI Guardrails, AI Observability, Model Evaluation, Hallucination Mitigation, Responsible AI, AI Governance

**AI / ML:** LangChain, LlamaIndex, LoRA/PEFT (QLoRA), Open-source LLM Deployment, WebGPU Inference, Fine-tuning, XGBoost, MLOps, MLflow, Anomaly Detection, MCP

**Big Data / ETL:** Apache Airflow, Spark, PySpark, Teradata, Informatica, Kafka, Huawei HetuEngine / DWS / MRS

**Backend / DB:** Spring Boot, Node.js, Flask, FastAPI, Express.js, Laravel, MySQL, Oracle, MongoDB, Redis

**DevOps / Cloud:** Kubernetes, Docker Swarm, Helm, CI/CD, AWS, GCP, Azure, Huawei CCE, n8n, Apache Superset, Metabase

## Education

---

**BSc Computer Science & Engineering: University of Asia Pacific, CGPA: 3.77/4.00**

Undergraduate research: Transfer Learning (model adaptation and fine-tuning)

## Open Source & Research

---

**HuggingFace:** [pesnik/pesnik](https://github.com/pesnik/pesnik) (1B fine-tuned model, 40 downloads) · [pesnik/pesnik-lora](https://github.com/pesnik/pesnik-lora) (QLoRA adapter, 26 downloads) · live GPU-based WebGPU inference at [pesnik.dev](https://pesnik.dev)

**PyPI:** 4 published packages ([pypi.org/user/pesnik](https://pypi.org/user/pesnik)): [superset-hetuengine-connector](https://pypi.org/project/superset-hetuengine-connector) (Superset → Huawei HetuEngine, first-of-kind) · [dagsonar](https://pypi.org/project/dagsonar) (Airflow task change visibility) · [diagnostics](https://pypi.org/project/diagnostics) (LLM ETL monitoring) · [tpt-builder](https://pypi.org/project/tpt-builder) (Teradata PT script builder)

**GitHub:** 200+ public repos ([github.com/pesnik](https://github.com/pesnik)); open PRs to Daytona (AI dev sandbox) and Excalidraw; DAGnostics (LLM ETL monitoring)

**Research:** ORCID: 0009-0006-4082-5527 · Kaggle: [pesnik](https://kaggle.com/pesnik)

## Awards & Achievements

---

**Professional:** Spot Award for AI solutions, CTIO Excellence Award (Airflow migration leadership), Game Changer Award (2×) — Banglalink Digital

**Academic:** Dean's Award (4×), Vice Chancellor's Award, 1st Place UAP Programming Contest, 1st Place National Essay Writing Competition (Red Cross BD), ACM ICPC Regional 36th (best) · 96th/2019 (UAP SpamBytes) · 252nd/2018 · 138th/2017 (UAP Override)