

M. Baizid Alam

Custody-grade governance discipline,
applied to **autonomous AI systems**
Banker · Engineer of Machines

Bankers can't build. Builders never answered to a regulator.

One person · six disciplines · nearly seventeen years.



OPERATOR · ID CONFIRMED

:50K: ORDERING CUSTOMER — CONTACT & COORDINATES			VERIFIED
EMAIL baizid.alam@gmail.com	CALL · WHATSAPP +88 0171 348 7996	HIRE · UPWORK upwork.com/freelancers	HIRE · FIVERR fiverr.com/s/NN7o0La
HUB aibony.github.io	LINKEDIN linkedin.com/in/mba2002	LIVE METRICS /stats.json	NINA UPDATES github.com/aibony/aibony ny

:32A: MESSAGE CONTENTS

:72:	Why This Niche — where the market pays	02
:20:	Profile · System Philosophy · Cognitive Profile	03
:23B:	NINA — Architecture & Honest Scoping	04
:32A:	What It Is Worth — the value ladder	05
:71A:	What Hiring Me Buys	06
:53A:	Independent Technical Due Diligence	07
:32A:	Core Competencies — eight modules	08
:50K:	Career Record — BASIC Bank PLC	09
:71A:	Education, Credentials & Leadership	10
:77B:	The Full Picture — strengths & weaknesses	11
:79:	How to Verify This	12

:72:

Why This Niche

The lane: custody-grade governance discipline, applied to autonomous AI systems in regulated finance — prompt-injection defence, scoped permissions, approval gates, runtime policy enforcement, audit trails, and secure fallback on constrained infrastructure. The market is paying for governed agents, not generic model-building. Public hiring and product evidence from Stripe, Brex, Adyen and Notion shows the highest-value AI roles inside fintech are about guardrails, auditability, policy enforcement, fraud and risk controls, and secure agent execution — not pure model-building. That is precisely the discipline a bank's SWIFT administrator, ISO 27001 lead auditor and trade-finance compliance owner already practises every day.

I design and implement governed autonomous agents for high-trust financial operations: scoped permissions, approval gates, audit trails, policy engines, runtime monitoring, and secure fallback on constrained infrastructure.

STRIPE

Inference-path security

An AI Security team builds tooling against prompt injection, jailbreaks and tool misuse; a separate SGRC role covers ISO 27001, PCI DSS, SOX and NIST control evidence.

BREX

Policy-enforcing agents

Agents enforce expense policy, auto-approve low-risk cases, escalate exceptions by risk level and monitor spend continuously — compliance experts become "workflow designers and policy engineers."

ADYEN

Agentic commerce control

Framed around explicit mandates, tokenization, SCA compliance, authentication and fraud prevention for merchant-controlled agent transactions.

NOTION

Zero-permission runtimes

Custom agents start with zero permissions, layered prompt-injection mitigations, runtime monitoring and mandatory confirmation for risky actions.

NIST · AI RISK MANAGEMENT FRAMEWORK

The regulator-facing vocabulary

NIST's AI RMF names prompt injection and agentic risk as documented concerns — so the language is no longer vendor marketing. It is the vocabulary an auditor and a risk committee will use. I already write in it.

WHAT THIS MEANS IN PRACTICE

Every one of those four companies is solving the same problem from the engineering side: how to let an agent act without losing the ability to explain, constrain and audit what it did. Banking solved that problem for money decades ago — message trails, segregated authority, mandatory stop rules, second-line review. I have spent nearly seventeen years operating inside those controls and the last months building an AI system that inherits them. The transfer is not theoretical; it is the architecture three pages on.

"Very few people combine SWIFT and ISO 27001 familiarity with autonomous AI architecture and local-first systems design in one profile. That intersection is a genuine moat."

— INDEPENDENT DEEP-PROFILE ASSESSMENT, 2026

:20:

Profile & System Philosophy

SWIFT ALLIANCE ACCESS ADMIN

ISO/IEC 27001:2022 LEAD AUDITOR

CBPR+ ISO 20022 TRAINER

AGM & DEPUTY HEAD, TRADE FINANCE

B.SC. MECH ENG · IUT

MBA SUMMA CUM LAUDE · NSU

Nearly seventeen years ago I entered a bank as a mechanical engineer who could not stop taking machines apart. Every posting since has been a new machine to master: branch operations taught me where risk hides in everyday transactions; credit taught me how money is actually analysed; trade finance taught me how documents move value across borders; SWIFT administration handed me custody of the bank's most critical infrastructure.

Mechanical engineers mechanise things. They take a process performed by hand and build a mechanism that performs it reliably, within tolerance, with failure modes understood in advance. I did not learn to code my way into AI — I mechanised DevOps. The substrate changed from steel to pipelines, agents and runtimes; the discipline did not change at all.

I specialise in local-first architectures under strict hardware and resource limits, and treat those limits as design inputs rather than excuses. What others solve with expensive cloud compute, I solve through structural optimisation, token discipline and deterministic governance.

"An investigative, systems-oriented builder who sees AI not just as a model interface but as infrastructure that must be governed, audited, and evolved over long time horizons — closer to a governance-first systems architect than a pure product hacker."

— INDEPENDENT ENGINEERING-PSYCHOLOGY PROFILE OF THE CODEBASE, 2026

:21:

Cognitive & Social Profile

Evidence from my career and NINA's architecture supports a profile of unusually fast technical abstraction and applied systems reasoning. I learn new tools rapidly, disassemble and troubleshoot machines without formal training, and administer complex banking and SWIFT systems by self-directed mastery. My mind searches for compression, leverage and waste reduction — one-click fixes, deterministic configuration, token efficiency, local-first execution.

Socially, I am a high-empathy systems thinker. Colleagues and superiors describe me as loveable, irreplaceable and overqualified; they come to me with both technical and personal problems because my advice is empathic and operational. That combination — high-bandwidth technical reasoning plus socially usable intelligence — defines me as a high-capability, high-trust node inside institutions. I am a scarcity-induced architect: I live under real hardware and budget limits and treat them as design inputs, not excuses.

THE OPERATING LOOP

RESEARCH → BUILD → REPEAT. NINA is the evidence that the loop works under real constraints. The next system built with it can be yours.

My meta-cognition is explicit: I watch how I think. Much of the day is spent in quiet, reflective problem-mapping, designing loops — audit, research, plan, implement, test, observe — that keep attention on long-horizon, high-leverage work. NINA is the externalisation of that inner loop: a machine that thinks about thinking, not merely about output.

:23B: NINA — Architecture (Proof, Not Product)


LOCAL-FIRST · MULTI-AGENT

2 GB VRAM CEILING

GOVERNED TRANSACTIONS

UNDER ACTIVE DEVELOPMENT

NINA is a personal, local-first autonomous AI infrastructure: a multi-agent orchestration system I designed, built and operate independently, built to run continuously on consumer hardware under strict resource constraints. Commands enter over Telegram, pass through a sanitisation gateway, and reach build and audit engines behind a governance layer — constitutional rules, compliance scoring, file registries, semantic memory. Inference is tiered and cost-aware, preferring local and low-cost paths before anything paid. Every change is a governed transaction rather than a blind edit.



Her active capabilities include offline semantic session caching, proactive system governance and auditing, AST proximity analysis and dependency mapping, continuous thermal and resource guarding, structured handoff capsules, and lossless context compression. She is evolving into a Strategic Operating System: an audited, deterministic local AI operating system with hardware and token autonomy, SWIFT-compliant modelling, and zero-hallucination causal reasoning. What I offer to employers and clients is the ability to design and coordinate similar systems for their context — not a downloadable repository.

THE GOVERNANCE IS NOT A DOCUMENT — IT IS A GATE

Most candidates for agent-security work can describe governance. Fewer have run a system where the governance actually refused something. NINA's approval gate has blocked merges that failed policy; component boundaries are enforced rather than advised; every state change leaves an audit record. That is the difference between a framework written for a slide and one that has said no to its own author.

SCOPED HONESTLY · AND MEASURED IN PUBLIC

NINA is **not a product, not for sale, and not production-grade**. She is a personal capability project under active development, started March 2026, with continuous commit history since 1 June 2026. Rather than print metrics that go stale the day they are typed, the system publishes them: commits, modules, test counts, lines of code and gate status are computed from the repository and re-synced every thirty minutes at aibony.github.io/stats.json, with a live architecture graph at architecture.html and a status page at systems.html.

:32A:

What It Is Worth

Every tier of what I deliver, from operational survival to competitive moat. The hierarchy is the argument: this profile operates at every level at once, not just the top or the bottom.

01 · SURVIVAL

Immediate ROI: reduce AI operating cost

Tiered inference — local models first, low-cost external second, premium providers only as fallback, re-evaluated on every call. The claim is simple: I reduce the inference bill from day one, and I can prove it from my own system's logs.

02 · SAFETY

Governed, deterministic AI — no surprises

Defined precedence rules. Differentiated trust levels per component. Explicit isolation of sensitive data. Bounded automated actions. Every change a governed transaction rather than a blind edit. Built for organisations that cannot afford a rogue agent.

03 · BELONGING

The bridge between engineering and governance

Most organisations have people who understand one side. The translation between "we built it" and "we can prove it is safe and auditable" is where projects die. That translation is the work.

04 · ESTEEM

Credibility with regulators and boards

ISO 27001 Lead Auditor. SWIFT administrator. Seventeen years of board-level memoranda. I give enterprises the language and the evidence to defend an AI deployment to auditors, central banks and risk committees — not just to engineers.

HOW THE LADDER IS USED

Most vendors sell at one tier. A cost tool sells survival; a consultancy sells esteem; a platform sells belonging. The argument of this page is that the tiers are not alternatives — an AI programme that saves money but cannot be defended to an auditor fails at tier two, and one that satisfies the risk committee while tripling the inference bill fails at tier one. I work at every tier because the compliance layer and the code layer are the same problem seen from two ends.

05 · SELF-ACTUALISATION

Competitive moat and intellectual legacy

Organisations working with me build intellectual property that cannot be bought off a shelf. Governed autonomous AI plus custody-grade compliance in one architecture is a moat most companies will spend years trying to acquire.

Read from the bottom, the ladder is a cost argument. Read from the top, it is a strategy argument. Read whole, it is the reason a single hire can matter more than a vendor contract: the person who reduces the inference bill and the person who defends the deployment to the risk committee are usually two different people, in two different budgets, who do not speak the same language.

:71A:

What Hiring Me Buys

NINA is the proof of concept. The product is a repeatable capability: **I research, build, repeat** — governed, auditable systems an organisation can both trust and defend. Six problems that show up in every serious AI programme, and what I do about each.

PROBLEM	WHAT I DO	HOW
Rising LLM cost	Tiered resource selection	Local and low-cost paths preferred before anything paid; the choice is re-evaluated continuously rather than configured once.
Vendor lock-in	Provider-independent architecture	Nineteen providers behind one abstraction with local-first fallback. No single point of commercial or technical failure.
Ungoverned AI	Governance as code	Precedence rules, differentiated trust levels and bounded actions, enforced at runtime rather than written in a policy nobody reads.
Data exposure	Explicit isolation	Sensitive data handled through a dedicated component under stricter isolation; local-first deployment for data-residency constraints.
Slow delivery	Structured, delegable specification	Work specified precisely enough to dispatch and verify. Sustained delivery cadence without growing headcount.
Fragile 24/7 AI	Layered resilience	Automated failure detection and recovery, thermal and resource guards, auto-restart supervision. Software plus physical safeguards.

SERVICE

AI cost engineering

Routing design and tiered inference for organisations watching the inference bill grow.

SERVICE

AI governance frameworks

Scoped permissions, approval gates, audit trails and policy engines — designed to survive an audit.

SERVICE

SWIFT / ISO 20022 consulting

From a practitioner who administers these systems today rather than having read about them.

Each row above is a problem I have met on my own system before meeting it on anyone else's — which is the point of building under constraint. The cost tiering exists because I could not afford the bill. The provider abstraction exists because a free tier disappeared overnight. The stop rule exists because an agent once wrote to a file it was not permitted to touch. None of it is theoretical; all of it was paid for once already.

:53A:

Independent Technical Due Diligence

SOURCE: EXTERNAL REPOSITORY REVIEW

REDACTED · CATEGORY-LEVEL ONLY

An independent technical review of the NINA repository found a local-first, continuously-running autonomous AI system built and maintained by a single owner-operator, using a development workflow that combines human-directed architecture with AI-assisted implementation — with sustained, high-frequency engineering activity over an extended period and a large, actively maintained backlog of fully specified work.

The system is designed as a personal autonomous operator rather than a conversational chatbot — intended to act inside real workflows rather than only respond to prompts. It runs on consumer-grade hardware as a continuously supervised background service configured to recover automatically from failure, with inference served through locally-hosted lightweight models and external providers behind a common internal abstraction layer.

ARCHITECTURAL CONCERN	TYPICAL CLOUD-AI APPROACH	THIS SYSTEM'S APPROACH
Compute location	Centralised, cloud-first	Local-first with cloud fallback
Resource selection	Static, configured once	Dynamic, continuously re-evaluated
Failure handling	Basic retry logic	Layered software and physical safeguards
Permission model	Often uniform access	Differentiated trust levels by component
Change governance	Automated tests, sometimes human review	Embedded priority policy plus explicit boundaries

The review named **governance-as-code** as the system's defining trait: a defined precedence among competing engineering priorities embedded directly into operating rules, with safety and data protection treated as non-negotiable constraints. Distinct execution components carry different trust and permission levels, and sensitive information is handled through a dedicated component under stricter isolation.

On engineering velocity, version history shows a sustained pattern of frequent, incremental changes merged in rapid succession during active development periods, alongside recurring automated maintenance keeping internal documentation in sync with the codebase. The project maintains an actively managed backlog in which each item is fully specified with defined scope, explicit boundaries, and a verifiable definition of done — the same discipline applied to work dispatched to human teams.

"Some foundational subsystems were still under active development at time of review. The system should be described as actively maturing, not fully production-hardened."

— TECHNICAL DUE DILIGENCE REPORT · REPRODUCED DELIBERATELY: I SELL HONEST ENGINEERING, INCLUDING ABOUT MY OWN WORK

I reproduce the unflattering finding deliberately. A governance practitioner who publishes only the favourable half of an external review has demonstrated, in the act of publishing, exactly the failure mode his services are supposed to prevent. The disclosed limitation is the credential.

:32A:

Core Competencies

MODULE 01 — TRADE FINANCE

SWIFT Administration

SWIFT operator since 2009, administrator since 2024. Alliance Access/Web administration and migration, CBPR+ ISO 20022 coordination and internal training, CSP compliance, sanctions screening, IRS FATCA and IDES Gateway, RMA and GPI tracker management, ASYCUDA, EDF, correspondent banking.

MODULE 02 — RISK / COMPLIANCE

Regulatory Risk & ISO 27001

Lead Auditor, ISO/IEC 27001:2022 (ISMS). Associate (AIBB, 2023) and Junior Associate (JAIBB, 2023) of the Institute of Bankers, Bangladesh. Basel III frameworks, Credit Risk Grading (ICRRS), AML/KYC, CTR/STR reporting, board-level regulatory and security memoranda.

MODULE 03 — BANKING BREADTH

Complete Domain Coverage

Nearly seventeen years across cash and general banking, credit (proposals, feasibility, recovery, write-off), foreign trade import operations, and head-office trade finance — teller counter to Board memorandum. Systems: SWIFT Alliance, Kastle Core Banking, TACPS, ASYCUDA, IDES.

MODULE 04 — AI INFRASTRUCTURE

Agentic Systems Design

Multi-agent orchestration, LLM integration, context compression, agent-driven code synthesis, structured delegation to AI agents — practised hands-on through NINA under strict hardware limits.

MODULE 05 — COST ENGINEERING

Multi-Source Orchestration

Deliberate design toward minimising ongoing inference cost through tiered resource selection — local and low-cost paths before paid services — with multiple inference sources behind a common abstraction and automated selection logic.

MODULE 06 — GOVERNANCE

Autonomous System Governance

Governance-as-code: defined precedence among competing priorities embedded in operating rules, differentiated trust levels per component, strict isolation of sensitive data, bounded automated actions, scoped permissions and approval gates.

MODULE 07 — RELIABILITY

Resilience Engineering

Automated failure detection and recovery, hardware-protective safeguards including continuous thermal and resource guarding, layered software-plus-physical resilience for sustained autonomous operation.

MODULE 08 — TECHNICAL LEADERSHIP

Structured Delegation

Fully specified, delegable task briefs with explicit boundaries and verifiable completion criteria; Board, regulator and foreign-counterparty correspondence.

THE PATTERN AN INDEPENDENT REVIEW NAMED: "CONSTRAINED AUTONOMY"

Four patterns recur throughout the codebase: gating new capability behind safety and hygiene checks; a single source of truth per component; bounded automated actions; and attached verification per subsystem. Every backlog item is written like a detailed contractor brief — defined scope, required context, explicit boundaries, verifiable completion criteria — a skill the review associates with technical leadership.

:50K:

Career Record — BASIC Bank PLC (2009 – Present)

Assistant General Manager & Deputy Head — Trade Finance Division

HEAD OFFICE · DEC 2024 – PRESENT · SOLE SWIFT ADMINISTRATOR, ENTIRE BANK

First state-owned bank in Bangladesh to complete the ISO 20022 CBPR+ migration — 100% outgoing traffic in MX format, one of only thirteen banks nationally at the time, per SWIFT's Regional Head. Trained every operator, supervisor and administrator personally through a Training of Trainers programme. Alliance Access and swift.com administration; CSP attested with 12 additional controls over the prior year; major release migration to v7.7.50 and transition onto new hardware with zero downtime and zero data loss; jump server implemented for access control. Restored the bank's participating FFI status under IRS FATCA via the IDES Gateway, then guided another state-owned bank through the same restoration as external technical consultant. Sanctions screening, authorisation of financial transactions to SWIFT, FI and correspondent banking, KYC/AML correspondence, remittance and cash incentive, Bangladesh Bank correspondence, ASYCUDA, data aggregation and Board memoranda. Introduced ISO 27001 governance patterns to change control, access management and auditability.

Senior Principal Officer — Trade Finance Division

HEAD OFFICE · SEP – NOV 2024 · FIRST SWIFT ADMINISTRATION CUSTODY

SWIFT and swift.com administration; patch discipline for Alliance Access and Alliance Web with zero disruption; RMA maintenance and GPI tracker updates; token management; KYC questionnaire; sanctions screening; Bangladesh Bank Online Monitoring System support; ASYCUDA and EDF operations; data aggregation and analysis. Where I learned that critical infrastructure rewards procedure and paranoia over improvisation.

Senior Principal Officer — Credit

MAIN BRANCH · NOV 2022 – AUG 2024 · SECOND MAN, CREDIT DEPARTMENT · BRANCH MANAGEMENT COMMITTEE

End-to-end credit proposals: new project evaluation and feasibility, working capital assessment, fresh loans, renewal, reschedule, interest waiver. Income projection from import and export operations; write-off implementation; cost of fund calculation; client financial statement review; ICRRS preparation; CL statements; recovery; Basel III analysis; wilful defaulter reporting; credit administration and audit compliance. Where I learned how a bank decides whom to trust, and what it costs when trust fails.

Executive Manager — Credit, Circle-4

HEAD OFFICE · APR – OCT 2022

Office notes for Head Office line management, Head Office Credit Committee and Board memoranda for credit proposals leading to approval; working capital assessment; analysis and accuracy validation of branch-submitted data; branch correspondence and monitoring. Where I learned what a Board actually reads, and how information must be shaped to survive the journey upward.

Manager — Foreign Trade (Import) In-Charge

MAIN BRANCH · JUL 2018 – MAR 2022 · STATE-INSTITUTION LC DESK · BRANCH CREDIT COMMITTEE

All import LC operations for Bangladesh's top state institutions — Bangladesh Bank, DGDP, CMSD, BPDB, GTCL — the branch's highest-value portfolio. Bond operation, credit report management, dealings with foreign suppliers' delegates, LC issue, document scrutiny, discrepancy raise and lodgement; student file operation, dividend remittance, D&B credit report realisation; online import monitoring, TM & IMP. Where I learned to hold the line on documentary discipline when the counterparty is a government.

Deputy Manager — Unit Head, Non Back-to-Back LC Unit

TRADE FINANCE DIVISION, HEAD OFFICE · AUG 2012 – JUL 2018

Six years leading a specialised LC unit: team supervision, preparation of proposals and statements, data extraction and aggregation, branch coordination, client information validation, Credit Committee and Board memoranda, SWIFT operations, RMA and correspondent banking. Where I learned to run a team by writing specifications clear enough that quality became repeatable — the same skill I now use to direct AI agents.

Assistant Manager

ASADGONJ BRANCH, CHITTAGONG · OCT 2009 – AUG 2012 · THE APPRENTICESHIP

SWIFT Alliance Messenger operation; DC issuance, amendment and cancellation; LC opening vouchers, commission and margin realisation; liability posting; acceptance verification; PAD creation and LTR adjustment; CIB inquiry; export bill collection, EXP certificates and proceeds realisation. General banking: outward and inward clearing, Nikash, TACPS, MICR cheque processing, OBC, account opening, CTR/STR, inward remittances. Cash operations. Temporary General Banking In-Charge and Cash In-Charge; deputations across Commercial Credit Division, Recovery Division, ICT Division (Head Office) and Jubilee Road Branch.

:71A:

Education

NORTH SOUTH UNIVERSITY · 2007 – 2009

MBA in Marketing — summa cum laude

CGPA 3.80 / 4.00. Awarded a merit-based 25% financial aid scholarship. NSU hosts the country's tier-one business school.

ISLAMIC UNIVERSITY OF TECHNOLOGY · 2003 – 2006

B.Sc. in Mechanical Engineering

CGPA 4.26 / 5.00 (converted 3.63 / 4.00). Full three-year scholarship from the Organisation of Islamic Cooperation. The foundation for systems thinking and engineering discipline.

NOTRE DAME COLLEGE, DHAKA · 2000 – 2002

Higher Secondary Certificate — 86.3%

One of the most selective and prestigious colleges in Bangladesh, and a tier-one institution by any national measure.

IDEAL SCHOOL & COLLEGE, MOTIJHEEL · to 2000

Secondary School Certificate — 86.0%

Among the top-ranked secondary institutions in the country. Tier-one schooling from the beginning, on merit throughout.

:70:

Credentials & Professional Training

CERTIFICATIONS

Formal credentials

Lead Auditor, ISO/IEC 27001:2022 — information security, cybersecurity and privacy protection (ISMS).

Associate, Institute of Bankers Bangladesh (AIBB), 2023.

Junior Associate (JAIBB), 2023.

CBPR+ ISO 20022 — coordinator and certified internal trainer.

CISSP — in progress.

Member, Institution of Engineers, Bangladesh (IEB).

TRAINING

Institutional programmes

SWIFT: Managing Alliance Web Platform, Operating Alliance Access, Managing Alliance Access (MTB Training Institute). SWIFT Administration and CBPR+ under ISO standard (Nelito Systems). International Trade Payment and Finance; UCP-600 and Related ICC Guidelines; Foundation Training on Banking (BIBM). Management of Non-Performing Loans; ICT Security & AML (BBTI).

:23D:

Earlier Roles & Leadership

2003 – 2009 · BEFORE BANKING

Teaching and apprenticeship

Senior Graduate Assistant to the Director, BBA Program, North South University (2008–2009); Graduate Assistant, MBA Program, School of Business, NSU (2007). Apprentice to the Managing Director, POWERMAN Engineering Ltd. Lecturer, OMECA Coaching Center (BUET admission coaching). Where the teaching instinct formed — the habit of making complex things learnable that now shapes CBPR+ training and agent specifications alike.

LEADERSHIP · EXTRACURRICULAR

Organising people

Sports Coordinator, MBA Club, North South University (2008–2009) — talent scouting, team organisation, performance monitoring. Student Coordinator, National Career Fair 2008 (NSU); Student Aide, 2007. Event Coordinator, Cultural Event Committee, Islamic University of Technology (2005–2006).

:77B:

The Full Picture

Honest scoping is the first skill of anyone selling governance. Strengths that are genuinely rare, weaknesses named rather than buried, opportunities that are time-sensitive, and threats already answered in the architecture.

STRENGTHS

A combination that is hard to assemble

Unique intersection. SWIFT administration, ISO 27001 lead audit and autonomous AI architecture in one person — this does not exist off a shelf.

Governance as runtime. Not a document. Enforced, auditable, versioned.

Constrained-environment mastery. A real commit history under a 2 GB VRAM ceiling proves resource discipline rarely found in cloud-native engineers.

WEAKNESSES · STATED BEFORE YOU ASK

Named plainly

NINA is not production-grade. A personal capability project under active development.

Solo operator. No team and no institutional backing; velocity depends on one person's evening hours.

Bandwidth. A full-time executive role limits me to structured, scoped engagements.

OPPORTUNITIES

Time-sensitive, and open now

Enterprise AI governance. Regulators increasingly require documented frameworks; few practitioners can build one from first principles.

Banking-sector adoption. Every bank is building AI. Almost none have someone who understands both the SWIFT network and the model layer.

Data-residency pressure. Local-first architecture expertise is scarce and getting scarcer.

THREATS · MITIGATED

Named, then answered

"Just use a cloud vendor." Local-first architecture proves vendor-independent capability; the moat is methodology, not model weights.

Credential inflation. Everyone claims AI expertise. Answered with a public dated repository, live metrics and independent assessments — not slides.

Role conflict. Answered with structured, scoped engagements and clear boundaries.

WHY THE WEAKNESSES ARE ON THE PAGE

Most profiles hide this quadrant. Mine prints it because the alternative is worse: a buyer who discovers a limitation after signing concludes it was concealed, and everything else becomes suspect by association. Named up front, the same facts are simply scope. I am a solo operator with a full-time executive role, so I take structured, bounded engagements — and that constraint produces better-specified work, not less of it.

Each threat above is answered by architecture rather than assertion. Vendor substitution is answered by a provider-independent design that already runs without any single vendor. Credential inflation is answered by a dated repository, a live metrics feed and assessments I did not write. Role conflict is answered by scope: bounded engagements with defined deliverables, which is how I have always worked inside a bank.

:79: How to Verify This

The repository is private. That limits what a stranger can independently check, so here is the honest split rather than an implied claim that everything is auditable.

VERIFIABLE RIGHT NOW — NO TRUST REQUIRED

Third-party sourced

The deployment clock on the site is fetched from GitHub's public API by *your* browser, not served from mine. The live architecture graph, the raw metrics feed, the two published essays and the GitHub profile behind them are all open to inspection.

SELF-REPORTED — TAKE ON TRUST, OR TEST ME

Generated from a private repository

Commit counts, module totals and gate results are computed by a script inside a private repo and published. You cannot audit them from outside, and I am not going to pretend otherwise. What I will do is walk any of it live in a screen-share, in whatever depth you want, and answer questions from the code rather than from a slide.

THE SAME DISCIPLINE APPLIES TO MY OWN CLAIMS AS TO THE SYSTEM'S

*A number you cannot check is worth exactly what you paid for it. I would rather name that limit than let it be discovered. Any bank can hire a SWIFT administrator; any firm can hire an AI enthusiast. What is hard to find is one person who holds custody-grade compliance discipline **and** builds governed AI systems with his own hands.*

:23E: CANONICAL LINKS — PERMANENT, CONTENT UPDATED IN PLACE			AIBONY.GITHUB.IO
HUB · DIGITAL CARD aibony.github.io	LONG-FORM PROFILE /profile.html	EVIDENCE / CASE FILE /evidence.html	LIVE ARCHITECTURE /architecture.html
LIVE SYSTEMS /systems.html	LIVE METRICS FEED /stats.json	CV · PRINTABLE /cv.pdf	CV · DARK EDITION /ocv.pdf
THIS DOCUMENT /portfolio.pdf	WRITING · MEDIUM medium.com/@aibony	WRITING · SUBSTACK aibony.substack.com	NINA UPDATES github.com/aibony/aibony

:57A: ACCOUNT WITH — REACH THE OPERATOR			OPEN TO ENGAGEMENTS
HIRE · UPWORK upwork.com/freelancers	HIRE · FIVERR fiverr.com/s/NN7o0la	EMAIL baizid.alam@gmail.com	CALL · WHATSAPP +88 0171 348 7996
LINKEDIN linkedin.com/in/mba2002	YOUTUBE aibony — all videos	DEMO VIDEO youtu.be/3cRpDgVsuMU	LOCATION Dhaka, Bangladesh

If a claim on this page matters to your decision, ask me to demonstrate it rather than describe it. I will open the repository in a screen-share and walk the code, the registries, the gate output and the failure log at whatever depth you want. That offer is the only thing on this page that costs me something, which is precisely why it is here.

Research → build → repeat. NINA is the evidence that the loop works under real constraints.
The next system built with it can be yours.