

# AZHAR KHAN

## AI Operations & Agentic Transformation

San Francisco, CA · 510-766-1727 · ap.apply@gmail.com · linkedin.com/in/azharpathan · azharorchestrates.com

**\$100M+ saved and 1M+ hours returned across a decade · 99% completion, 26M runs, one platform · under 3¢ per business outcome**

### SUMMARY

Leads enterprise AI transformations from discovery through production operations. Walks into the operation, maps the process as it actually runs, redesigns the operating model so AI agents handle the repeatable work and people handle judgment, makes the platform decisions, hires and directs the team that delivers it, and runs it in production. Three enterprises run on it today: a top-5 US fintech (an agent assembles the fraud case, one analyst decides), a national home-care platform (16 jobs in 10,000 reach a person), and a Microsoft-published utility operation (one invoice in 100 does). Set the strategy, delivered each, and directs the teams that run them. Each is a standing operation reporting to the CTO or COO.

### EXPERIENCE

#### Concentrix

2019 - Present

Associate Director, Agentic AI Engineering & Automation

San Francisco, CA

Grew the agentic automation Center of Excellence from 4 people to 30+: \$10M+ delivery revenue under management, architects, engineers, analysts, and delivery managers across the US, Mexico, and India, two to three enterprise transformations running at once. Forward-deployed inside each client's operation. Every benefit case clears a centralized CoE council. Five turnarounds across five industries.

- ◆ **BotScope, the control plane:** led product development of the fleet observability and governance layer (tracing, step-level failure attribution, MTTR, outage detection, cross-bot telemetry), 2019. Adopted across the practice; runs every engagement. Recognized by Forrester, 2021. Now carries agent identity, per-action entitlements, a live execution graph, and token economics.
- ◆ **CoE-in-a-Box:** centralized hub managing the full use-case lifecycle from discovery through production operations and FinOps; deployed across all five enterprise programs. A five-day sprint replaces multi-month assessments; the first phase ships at no new spend.

**RECOGNITION** · Microsoft-published case study (utility invoice processing) · Forrester recognition (BotScope, 2021) · UiPath AI10 Award

### CORE COMPETENCIES

**AI Operations & Governance:** AI Governance in Production · Decision-Rights Architecture · Evals & Observability · Drift & Hallucination Monitoring · AI FinOps (Cost per Outcome) · LLMOps / AgentOps · Multi-Agent Orchestration · Context Engineering & RAG · Human-in-the-Loop Design

**Discovery & Operating Model:** Process Discovery & Mining (Celonis) · Operating-Model Redesign · Digital-Twin Modeling · Platform Selection & Build-vs-Buy

**Platforms:** Vertex AI · Azure OpenAI · Azure AI Document Intelligence · LangGraph · BigQuery · Decagon · Sierra · Workato · Agent Studio · UiPath · Appian · Datadog · Glean · Braintrust · Power Platform · Salesforce · Workday · NetSuite

### PRODUCTION OPERATIONS

Client operations delivered and run under the Concentrix role.

#### Top-5 US Fintech · Agentic Fraud and Risk Operations · 2025 - Present

660K+ fraud cases a month · 7 lines of business · ~\$12M annual benefit · a human on every regulated decision

Seven financial-crimes-compliance programs. Baseline: each case built by hand across six to eight systems through a 21-step review of 100 to 150 data points, wrong a third of the time. Set the strategy; directs the team running the platform in production across Asia and LATAM. Global rollout commissioned; EMEA in design under the EU AI Act.

- ◆ **Platform selection, run as the buyer:** evaluated Decagon against Sierra on live case volume and deployed each where it won. Both in production across the seven lines of business.
- ◆ **Redesigned the work first:** atomic-step breakdown, disputes re-sequenced so denial parameters evaluate first, a 12-to-15-screen investigation collapsed into one workspace. Handle time down a third; CSAT up 13 points.
- ◆ **Agent builds, human decides:** an Automated Due Diligence agent pulls 129 data points from every system a case touches, scores them against the customer's baseline, checks SAR thresholds, and hands the analyst a decision-ready file with the report drafted. Every case auditable end to end.
- ◆ **Architecture:** orchestrator over specialized sub-agents (ingestion, investigation, reporting) on shared state; gradient-boosted risk scoring, tiers 1 to 4; hybrid vector and keyword retrieval over dispute history, chat logs, and ledger notes.
- ◆ **Controls:** evals on curated golden datasets for hallucination, retrieval precision, and compliance drift; step-level tracing; guardrails for prompt injection and SAR thresholds; PII tokenized and access-scoped end to end; prompt caching, token-cost limits, tool-call rate limits.
- ◆ **Built once, used seven times:** one spine, three checks (denylist, event-stream, UAR/SAR), serving account takeover, disputes, unauthorized transactions, elder financial abuse, identity verification, legal name change, and card signature. Each region deploys off the last.

#### Global Bank · Fraud & Risk Operations Transformation · 2024 - 2025

Eight fraud journeys, application fraud to scams and card disputes · 188-FTE discovery baseline across 40 systems · designed tiered agentic model for 2.48M fraud calls a year

Five-day discovery sprint across 188 FTE, 40 systems, and eight fraud journeys (application and first-party fraud, scams, card disputes, fraud admin on SLA stop-clock mechanics). Redesigned the operating model across four workstreams; 62 journey-system touchpoints collapsed to 3 intervention points, evidence chain intact. Directed from discovery through operating-model design; now in production. Re-engaged by the bank for its fraud and complaints roadmap.

- ◆ **Case journey:** 13 steps on 6 systems to 5 steps on 1; roughly 20 applications collapsed into one case view; junior operators work cases that needed specialists. Auto-allocation, downstream write-back, automated letters, GenAI case summaries, twenty-plus post-decision actions per outcome (CIFAS loads, fund-securing, bureau deletions). Intelligent Operations CoE stood up against a \$3.29M five-year savings plan.
- ◆ **Tiered model for 2.48M fraud calls a year:** automated resolution at level one, agent-assist at level two, human investigators on romance scams, high-loss, and vulnerable-customer cases.

## National In-Home Care Enterprise · Core Operating Platform, Governed Data & Agent Fleet · 2021 - Present

\$30M+ impact · 11 business functions on one governed platform · 99% completion across 26M runs over five years · 16 jobs in 10,000 reach a person

Baseline: point-to-point manual handoffs across 45+ disconnected applications. Designed the operating platform on Workato ONE in 2021 and has directed it since: 176M tasks to date, 12 US regions, reporting to the CTO, shipping against quarterly release cycles. Twenty AI use cases now run on top, no rebuild underneath.

- ◆ **Workforce:** 10,000+ hires a year across recruiting, background checks, HRIS, and payroll, zero orphaned accounts; 23 lifecycle scenarios absorbed an 18% hiring surge without a rebuild. Retention on the same platform: rewards, earned-wage access, auto-generated caregiver bios.
- ◆ **Revenue:** \$8.9M billed a week; annual rate change run as one orchestrated event across the entire client book; payments auto-matched to invoices; conversions written back to ad platforms; month-end close down 40%.
- ◆ **Finance:** multi-currency across subsidiaries; refunds cleared form to payment unattended; contractor invoices generated from delivered care; suppliers synced Workday to Coupa. One golden record per caregiver, client, and location; \$1.33M in leakage closed. Records decoupled from source systems, swappable without touching the process.
- ◆ **20 AI use cases in production** on the same platform, no rewrites: routed per use case across Vertex AI, Agent Studio, and Workato by business-function preference. An agentic recruiting pipeline replaced two hiring vendors. Decision rights per action; an enterprise MCP layer segregated by control, access, and use case; CCPA-aligned; every release through eval gates, drift monitoring, and a reconstructable audit trail.
- ◆ **Operations:** ITSM triage agent routes ~90% of a 30,000-ticket queue with no person in the path. Five years, no business-interrupting outages, through nine acquisitions and three sponsor changes. Per-action identity, 132 monitored tables, fail-closed recovery, design gates before build. Under 3 cents an outcome; consumption down 39% year over year as volume grew.
- ◆ **Readiness gate:** scored 16 capabilities of a separately bought AI estate, ruled the nominated use cases unsafe to launch, shipped a phased sequence, Phase 1 at no new spend. Roughly 40% of nominated use cases fail the review.

## Global Utility-Bill Management Operation · Data Platform & Energy-Intelligence Transformation · Ongoing

700K invoices a month toward 1M · \$2.5B in payments · 180+ clients · 20,000+ vendors · Microsoft-published case study

Baseline: a 700-person back office running the utility-bill lifecycle by hand across four groups and three billing models on a 17-year-old platform; a prior automation broke because the system lacked a unified view of state to validate against; half of extracted volume hit exceptions, four in five QA flags fired false. Directed the turnaround: diagnosed the data problem first, redesigned the operating model against a digital twin, and runs it in production. Forecasting, rate optimization, benchmarking, and Scope 1 and 2 reporting run off the reconciled data.

- ◆ **One reconciled record** of spend, usage, demand, and rate across 20,000+ vendors and 300+ layouts, through a three-level charge dictionary. The twin runs as a pre-exception validation layer, catching deviations before they fire and naming the next automation.
- ◆ **Capture accuracy 65% to 99%:** three-path confidence routing (OCR, trained models, generative extraction) across multi-language, multi-page invoices to meter and charge-line level. Handle time 8 minutes to under 1; cost per invoice \$0.58 to \$0.21, roughly \$3M a year; one invoice in 100 reaches a person. No automation goes live without a baselined process model.
- ◆ **Onboarding 90 days to 30:** all ten phases re-engineered, the two on the critical path automated; account templates pull 149 fields from invoice PDFs; the AP funding file builds in two to three days. Each idle month held roughly \$800K of contracted revenue.
- ◆ **Funding and reconciliation:** a \$50M jump in monthly card volume absorbed with no added headcount, ~1,900 payments cleared a day; historical loads gated at 99.5% reconciliation; remittance-address audit down to three people.
- ◆ **Team:** 290 document keyers reshaped into 65 pattern-builders and exception analysts across nine specialized teams; 225 roles freed for new clients and the climb from 700K toward 1M invoices a month. Agent-callable through the same MCP pattern; extended to freight, manifests, and leases by configuration.

## EARLIER EXPERIENCE

### UiPath

2018 - 2019

*Professional Services, Solution Engineering*

- ◆ **BP and BPX Energy:** delivered 70+ automations worth \$20M, eliminated 200K hours, and owned the M&A and SAP S/4HANA automation roadmap.

### ProKarma

2016 - 2018

*Automation Consultant*

- ◆ **US telecom carrier:** built the enterprise automation governance, intake, and delivery frameworks; invoicing errors fell 85% and processing time fell 90%.

## FOUNDATION: A DECADE OF PRODUCTION AUTOMATION AT ENTERPRISE SCALE

Concentrix Catalyst (formerly PK) and earlier roles: a decade of enterprise automation programs.

- ◆ **Fortune 200 insurer, claims:** AI-driven auditing on 540K claims a month; \$7.4M recovered at a 6.5x return where a 250-person manual operation had capped out.
- ◆ **Fortune 200 insurer, commissions:** 1M invoices a year for 12,500 clients on a three-person team. UiPath AI10 Award.
- ◆ **Fortune 200 insurance services, finance operations:** general-ledger process automated; mean time to post the GL cut to 8 minutes; days sales outstanding down a week.
- ◆ **Fortune 500 manufacturer:** enterprise automation program across four countries, \$6.5M+ returned, AP rebate cycle 90 days to 2. Generative-AI pipeline (Stable Diffusion, ControlNet, LangGraph) turned engineering drawings into catalog imagery for 13,000+ products, weeks to hours; Concentrix case study.
- ◆ **Top US telecom:** 250-automation program, \$18M returned; one redesign alone returned \$988K a year and freed 33 roles.

## EDUCATION

M.S. Computer Science, San Francisco Bay University · B.Tech Computer Science, Vardhaman College of Engineering