Chaitrishodaya

COMPANY PROFILE & CATALOGUE · 2026

PRODUCTION AI THAT PASSES THE AUDIT, NOT JUST THE DEMO.

We design, build and govern AI systems for banks, NBFCs and enterprises, with evaluation gates, adversarial testing, secure retrieval and tamper-evident audit trails built in from day one.

WHO WE ARE

THREE WAYS WE HELP

Every engagement starts from what you need to achieve, not from a list of tools.

PROVE YOUR AI WORKS

YOUR SCORES SAY 94%. IS THAT TRUE?

Golden datasets, adversarial tests and statistical checks that show how your AI really behaves.

GOVERN YOUR AI

CAN YOU PROVE WHAT YOUR AI DID, AND WHY?

Access control before retrieval, data masked by role, answers with sources and a tamper-evident audit trail.

RUN YOUR AI SAFELY

WOULD YOU KNOW IF QUALITY DROPPED TODAY?

Release gates, live monitoring, fast rollback with human sign-off, and cost routing.

THE FOUNDERS

FOUNDED BY BUILDERS

KS

FOUNDER · AI & ENTERPRISE ARCHITECT

B KUMARA SWAMY

13+ years in enterprise and software engineering. Architected an autonomous AI automation platform that cut manual regression by 70% across 2,350+ CI builds, and designed and built eleven AI systems end to end.

MJ

CO-FOUNDER · ENTERPRISE ARCHITECT

DR MANOJ KUMAR JAIN

20+ years across IBM, PwC, Oracle, TCS, HCL and Wipro. Leads transformation from business vision and strategy through architecture, engineering and AI to operations and value.

Measured, not claimed
Every number traces to a real system.

Honest about stage
Pilots are called pilots; research is labelled.

People stay in charge
AI investigates, humans decide.

SYSTEMS

ELEVEN SYSTEMS, EACH ONE PROVEN

AI QUALITY & TRUST VERIFICATION

MTS Metric Trust Score

CORE · 3-CHECK VERIFIED

PROBLEM
A dashboard shows 94% whether it rests on 12 data points or 2,000.

WHAT IT DOES
Checks every AI metric for too little data, outliers and drift before anyone trusts it.

RESULT
Three independent checks: modified Z-score, IQR cross-check and EWMA completeness.

AIPQ Prompt Quality CI Gate

CORE · AUTO-REVERT

PROBLEM
A prompt edit can quietly make answers worse, and nobody notices until users do.

WHAT IT DOES
Versions every prompt like code and scores each change against a golden dataset.

RESULT
A change that dropped quality to 0.60 was reverted to 0.93 automatically, in four minutes.

BCT Behavioral Contract Testing

CORE · ADVERSARIAL-TESTED

PROBLEM
Teams set rules for their AI but never learn how much pressure breaks them.

WHAT IT DOES
Attacks each rule with step-by-step adversarial pressure and reports exactly where it fails.

RESULT
An installable package (pip install bct-framework) that any project can run.

MULTI-AGENT ORCHESTRATION

QAIP QA Intelligent Platform

LIVE · 84% COST CUT

PROBLEM
Writing tests by hand cannot keep up with how fast software changes.

WHAT IT DOES
Reads a GitHub repo, scores the risk and generates runnable Playwright tests, learning each sprint.

RESULT
Cost per run cut from ₹50 to ₹8, an 84% reduction, with quality held.

ZENTRAVIX Organisation Intelligence Platform

LIVE · 3-SYSTEM RAG

PROBLEM
Every layer of a company answers questions from a different, outdated slice of data.

WHAT IT DOES
One intelligence layer across three systems, with answers tailored to each role from CEO to developer.

RESULT
Three systems unified, with RAG answers from live organization data.

PREDICTIVE RISK & MONITORING

AIMO — AI Incident Monitoring

CORE · 5 FAILURE TYPES

PROBLEM
AI can get worse while every server dashboard still shows green.

WHAT IT DOES
Watches for five AI-specific failure types in parallel, including in its own alerting.

RESULT
It caught a real bug in its own alerting logic during validation.

SCIP — Supply Chain Intelligence Platform

LIVE · 80+ ENDPOINTS

PROBLEM
Suppliers that look independent can share one hidden upstream risk.

WHAT IT DOES
A full supply chain platform with cost workflows and explainable supplier risk scoring.

RESULT
80+ endpoints, and a critical (P0) bug caught before launch.

TEST AUTOMATION AT SCALE

Universal Test Framework — One Test Flow, Five Automation Tools

LIVE · 5 FRAMEWORKS

PROBLEM
Teams copy the same test into several tools and maintain every copy.

WHAT IT DOES
Describe a test once and generate Playwright, Selenium, Cypress, TestNG or Cucumber code.

RESULT
Five working adapters, and it found a real bug in a separate application.

Locator Framework — Playwright Modular Locator Framework

CORE · 9/9 PASSING

PROBLEM
UI tests break when one element locator changes, even if the feature works.

WHAT IT DOES
Falls back to the next locator strategy and logs exactly which one healed the test.

RESULT
9 of 9 tests passing, including against a live website.

ADAPTIVE AI LEARNING SYSTEMS

ARIA — Socratic AI Tutor, 35 languages

LIVE · 100% COMPLIANT

PROBLEM
AI tutors give in and hand out answers when students push hard enough.

WHAT IT DOES
A Socratic tutor in 35 languages that guides with questions and holds under pressure.

RESULT
Real compliance raised from 22.2% (despite a 94% evaluation score) to 100%.

SECURE AI FOR REGULATED DATA

Banking Document Intelligence — Secure RAG for Regulated Data

PILOT · ACCESS BEFORE RETRIEVAL

PROBLEM
Typical document AI searches everything first and hides data afterwards, when it is already too late.

WHAT IT DOES
Answers staff questions from bank documents, checking access before any document is searched.

RESULT
Access before retrieval, data masked by role, a tamper-evident audit trail and local models.

SERVICES

FROM FIRST ARCHITECTURE SKETCH TO PRODUCTION

01 **AI System Architecture & Development**

We design and build complete AI-powered platforms from the ground up.

02 **Enterprise AI Integration**

We connect AI into your existing systems and workflows, unifying data and intelligence across teams.

03 **AI Quality & Reliability Engineering**

Evaluation, monitoring, and adversarial testing layers that keep production AI systems trustworthy.

04 **AI Cost & Performance Optimization**

Smart model routing and pipeline efficiency that cut cost without cutting quality.

05 **AI Security & Compliance**

AI systems fail in ways traditional software doesn't.

HOW WE WORK

1 · AUDIT

Trace the pipeline and find where evaluation and monitoring are missing.

2 · DESIGN

The specific gates your system needs, for your real failure modes.

3 · BUILD

Integrated into your existing CI/CD and monitoring stack.

4 · VERIFY

Stress-tested with real adversarial cases before it is called done.

START WITH A FREE 30-MINUTE REVIEW

Tell us about one AI system. We reply within one business day, from a founder.

chaitrishodaya.com linkedin.com/in/kumara-swamy-7731b020