

Enterprise AI Lifecycle Platform

Discover → Evaluate/Prioritize → Recommend AI Approach → Recommend right LLM/SLM → Build Agent → Execute → Govern → Measure

Avagama is Avaali's Enterprise AI Lifecycle Platform.

Avaali is an enterprise technology firm with a 13-year delivery track record across source-to-pay, AI, automation, master data and information governance serving large enterprises across India, Middle East and South-East Asia.

300+

large enterprises
served

40–70%

cycle-time
reductions delivered

50,000+

suppliers enabled
through Velocious

15+

years enterprise
delivery



Velocious – Esourcing and P2P solution



Available on SAP Store



ISO 27001:2013 Certified



Partnerships: SAP · OpenText · UiPath · Microsoft Power Automate ·
Prospecta

Trusted by 300+ Enterprises Across Industries.

Manufacturing • Pharma • Healthcare • Energy • PSU

Automotive • Retail • BFSI • Media & Conglomerates

Manufacturing and Industrial



Pharmaceuticals & Healthcare



PSU, Energy & Gas



Automotive



Retail & Consumer Goods



BFSI, Media and Conglomerates



"The common thread across every engagement: Avaali improves operating performance where processes cross systems, suppliers, documents, data and decisions."

What Our Customers Say

“

We're thrilled with Avaali's expertise in automating our accounts payable process. The collaboration reflects our commitment to leveraging cutting-edge technology, and we're already seeing significant gains in agility, visibility, and governance across our AP operations.

Mandar Thakur Desai

Deputy CIO – IT & ERP

Tata Projects

“

Digital transformation isn't just about technology - it's about reimagining how we operate for the future. Velocious has fundamentally elevated our procurement capabilities, delivering the operational agility and intelligence we need to stay ahead.

Prem Kiran Udayavarma

CIO

Aditya Birla Renewables

“

Managing SAP-linked documents across a vast operational footprint was becoming increasingly complex. Avaali's digital workplace expertise helped us establish a centralized, compliant, and efficient way to manage content across the enterprise.

Sanjay J. Joshi

Chief Technology Officer

Parle Products Private Limited

Boards Are Funding AI. Enterprises Are Stalling.

Three decisions are killing enterprise AI programmes before they start.

1

Which Use Cases Actually Matter?

Without structured evaluation, the highest-paid opinion in the room decides where AI goes. High-value opportunities are missed. Low-value experiments consume budget and credibility.

2

Which AI Model Is Right for Each Process?

GPT-4o, Phi-3, Mistral, Claude — every process has a different optimal model. Most enterprises choose once and apply everywhere. That's expensive and often wrong.

3

Where Is the Proof of Return?

Without metrics defined before build, AI investments become unaccountable. Boards stop funding programmes they cannot measure. CFOs write off "innovation spend".

11% of technology leaders say they are fully prepared for AI deployment at scale — yet most programmes stall before production. (McKinsey, 2025)

74% of enterprises expect AI-driven revenue growth. Only 20% deliver it. The gap is not ambition — it is execution discipline. (WEF, 2025)

We Looked at Every Major Platform in the World.

None of Them Do This.

Process Mining Platforms ✓ Strong at discovering and mapping processes across the enterprise	MISSING: Stops at analytics. No LLM recommendation, no agent build, no execution. Discovery without action.
ERP & Suite-Embedded AI ✓ AI built into existing enterprise systems - fast to activate within the suite	MISSING: Platform-locked to one ERP or workflow suite. Cannot evaluate processes outside the system. No model agnosticism.
Ecosystem Agent Builders ✓ Powerful agent creation within established cloud and CRM ecosystems	MISSING: No process discovery. No model recommendation. Agents live only inside one vendor ecosystem - no cross-platform intelligence.
DIY on Cloud AI Platforms ✓ Maximum model flexibility and cloud-native scalability	MISSING: No discovery framework. No blueprint library. No recommendation logic. Engineering teams build from scratch - months before the first production agent runs.

Avagama is the only platform that connects the entire chain: Discovery → Evaluation → LLM/SLM Recommendation → Agent Build → Autonomous Execution.

One Platform. The Complete Chain.

No other platform connects all six steps in a single integrated intelligence layer.

1**DISCOVER**

Process Discovery

Upload documents or search by function and industry. Avagama surfaces use cases your team would never identify manually.

2**EVALUATE**

Multi-Dimensional Evaluation

Each use case scored across decision complexity, data readiness, strategic fit, risk and compliance — zero subjectivity.

3**RECOMMEND**

LLM / SLM Recommendation

The only platform that analyses your process and recommends the optimal AI model — GPT-4o, Phi-3, Mistral, or your own - before you build a single agent.

4**BUILD**

Agent Auto-Build

Build agents with pre-defined blueprints, multi-framework support, and full lifecycle management. 31 blueprints across 17 industries - ready in hours.

5**EXECUTE**

Autonomous Execution

Agents act - not just recommend. With full audit trails, cost tracking, and observability built in from day one.

6**GOVERNANCE**

Continuous Governance

Policy enforcement, human-in-the-loop approvals, and compliance monitoring applied across every agent, every step of the chain.

Three Things No Competitor Does.

01

LLM / SLM Recommendation Per Process

The capability no global platform offers.

- ❑ Every other platform lets you bring a model.
- ❑ Avagama analyses the specific characteristics of your process - volume, complexity, exception logic, compliance sensitivity, latency - and recommends the optimal model before you build.
- ❑ This decision currently lives in your head, or is made by guesswork. Avagama makes it objective and auditable.

02

Discovery to Execution — One Integrated Platform

The missing layer between process mining and agentic AI.

- ❑ Process mining stops at insights, not execution. ERP AI stays confined to one suite. Agent builders require you to know what to build, while DIY cloud platforms offer flexibility but no blueprints, governance, or fast path to production.
- ❑ Avagama is the only platform where a process enters one end and a governed, production-grade agent exits the other - every decision documented, every model choice justified, every outcome measurable.

03

31 Process Blueprints Across 17 Industries — Ready Now

Not templates. Pre-validated intelligence.

- ❑ 31 production-tested blueprints define the right AI approach for the right process across finance, procurement, healthcare, manufacturing, logistics, legal, and more.
- ❑ Customers can launch their first AI agent in hours, backed by a blueprint library that grows with every deployment.

The Capability Map No Competitor Can Match.

Capability	Avagama	Process Mining Platforms	ERP & Suite Embedded AI	Ecosystem Agent Builders	DIY Cloud AI Platforms	Point Solutions Stitched Together
Process Discovery & Evaluation	✓	✓	🕒	✗	✗	🕒
LLM / SLM Recommendation per Process	✓	✗	✗	✗	✗	✗
Discovery → Execution (single platform)	✓	✗	✗	✗	✗	✗
Autonomous Agent Execution	✓	✗	✓	✓	✓	🕒
Bring Your Own LLM (truly agnostic)	✓	✗	✗	✗	✓	🕒
31+ Industry Process Blueprints	✓	✗	🕒	🕒	✗	🕒
Enterprise Governance & Audit Built In	✓	🕒	✓	🕒	🕒	🕒
Agent Auto-Build from Blueprints <i>(Available in Beta)</i>	✓	✗	🕒	✓	✗	🕒

✓ = Full capability 🕒 = Partial / requires add-on ✗ = Not available

A Four-Step Engine From Process Input to Prioritized Action.

1

Seamless Process Input

Upload existing process documents, or search by business function and industry. Avagama surfaces use cases your team would never have identified manually - including those buried in exception logs and workarounds.

2

Multi-Dimensional Evaluation

Each use case is scored across decision complexity, data readiness, strategic fit, risk and compliance sensitivity. No subjectivity. No politics. A ranked, defensible pipeline your board can interrogate.

3

LLM / SLM Fitment Analysis

A unique AI Fitment Score recommends the optimal model and execution path for each process: RPA, Augmented AI, or Agentic AI - and specifically which model to use. No guesswork. No expensive mistakes.

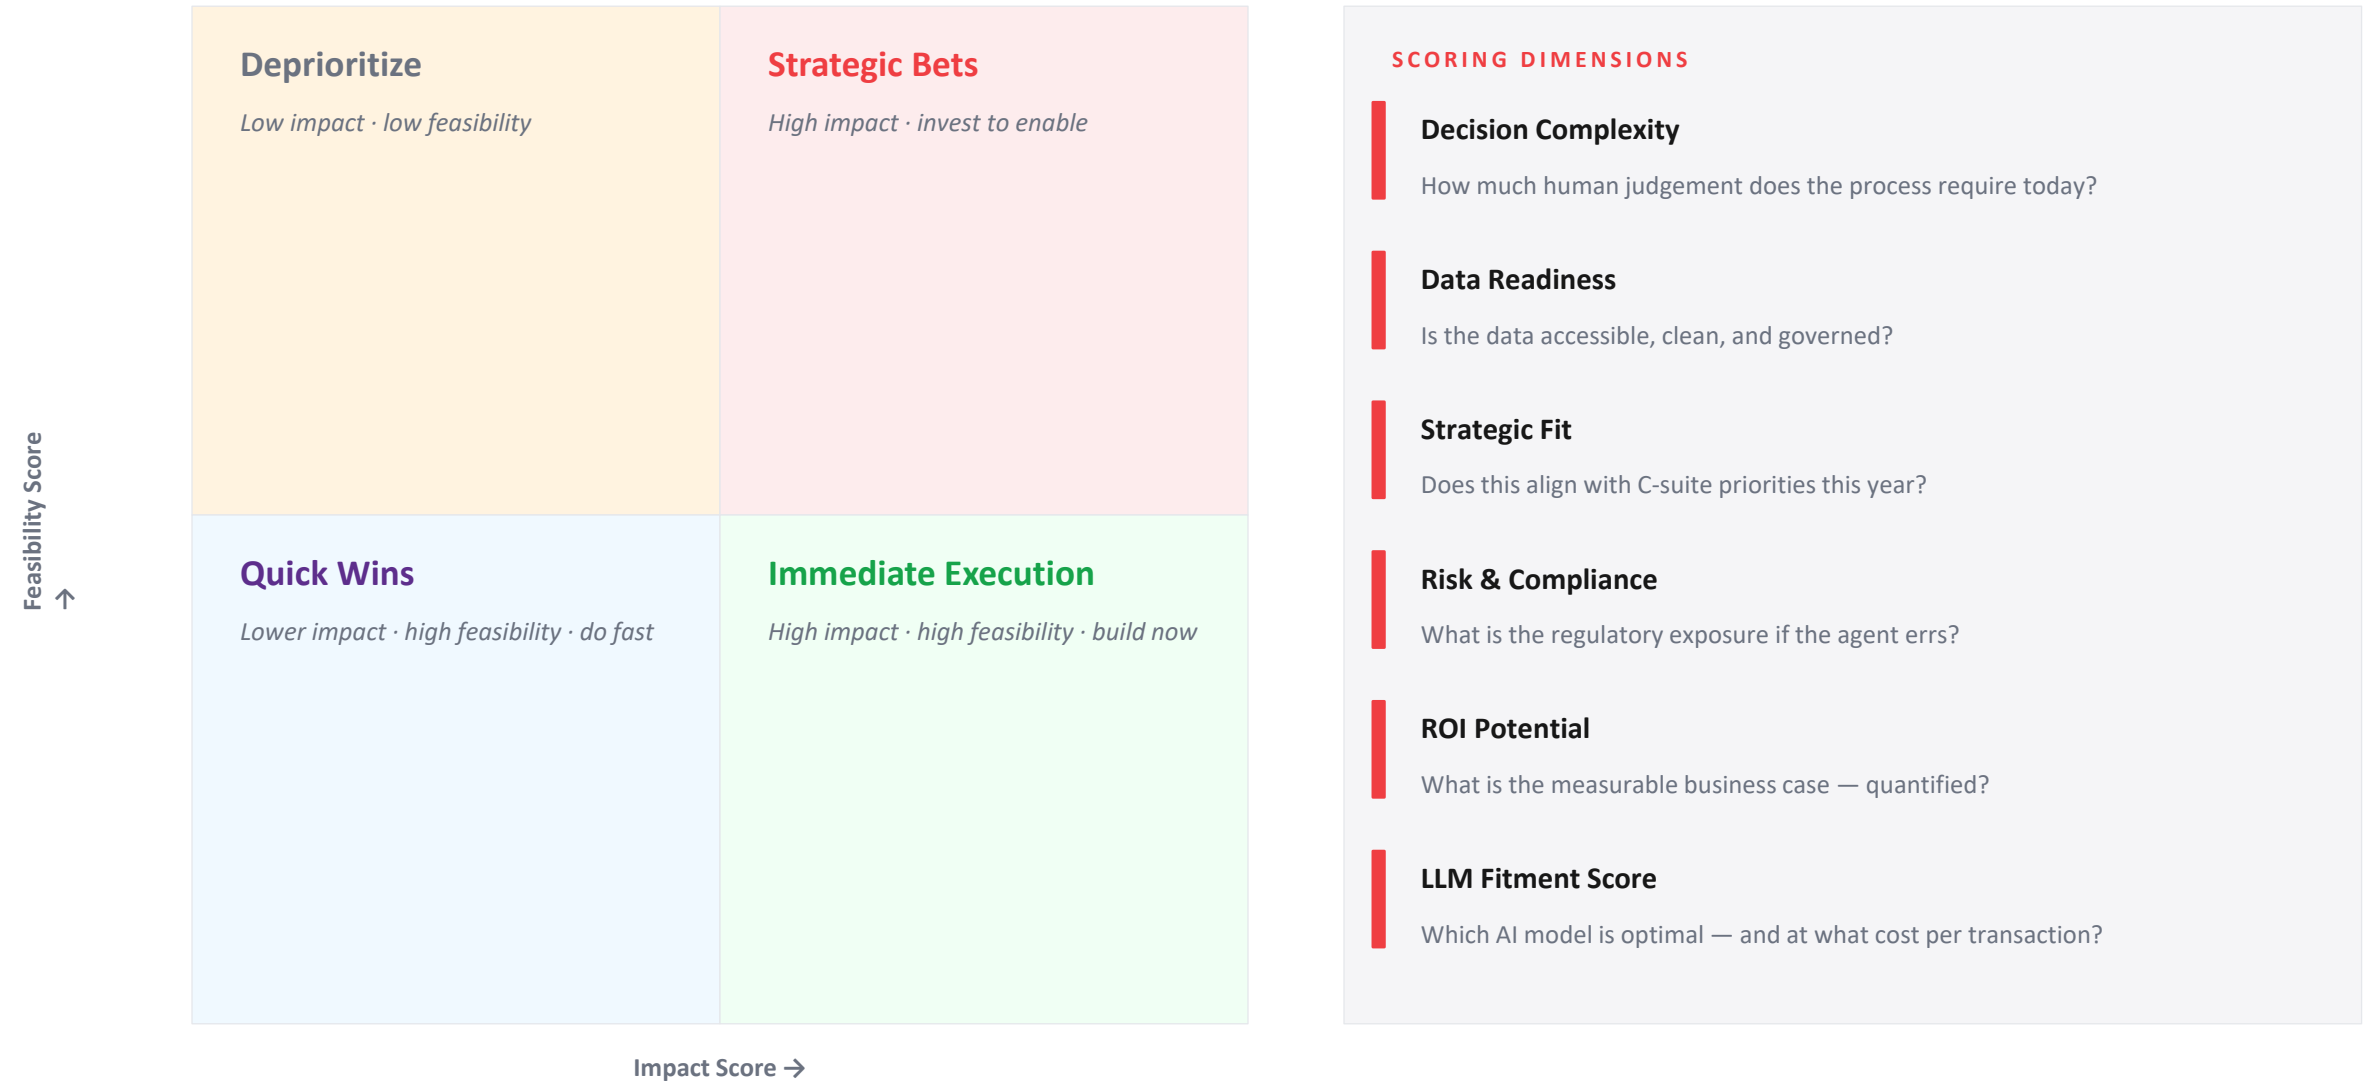
4

Strategic Use Case Ranking

A board-ready, prioritized pipeline. Your team knows exactly what to build first, why, which AI powers it, and what return to expect - before a single line of code is written.

From Gut-Feel to Board-Ready Investment Decisions.

Every use case is scored across impact and feasibility. Your pipeline becomes objective, auditable and defensible.



Build, Deploy and Execute Agents — Your Way.

Every major agent framework. Enterprise orchestration. Multi-agent workflows. Production-grade from day one.



AI Agent Creation

Build and deploy agents with persona design, tool integration, memory management, and end-to-end lifecycle control across all major frameworks.



Multi-Agent Workflow Orchestration

Stateful, multi-agent workflows with execution tracing and observability from day one. Coordinate agents on complex, multi-step enterprise processes.



Execute Agent Functionality

Agents don't just recommend - they act. Auto-execution with configurable confidence thresholds, human-in-the-loop escalation, and full audit trails for accountability.



Continuous Agent Enhancement

Ongoing upgrades to agent reasoning, memory, tool integration and output quality - delivered without disruption to live agents running in production.



Structured Evaluation (LLM-as-Judge)

Automated, structured evaluation of every agent output. Quality assurance that scales with usage - so compliance and finance teams have defensible evidence, not hope.



AI Observability & Cost Transparency

Full visibility into every agent interaction, decision trace, and token cost. Per-user, per-module cost analytics - the dashboard your CFO has been asking for.

Agents Work Where Your Teams Already Work.

No new interface to adopt. No workflow disruption. No change management project.

Slack Connector

Deploy agents directly into Slack for conversational workflows. Approvals, alerts, queries - all handled where your teams already communicate.

Gmail Connector

Let agents read, draft, and act on email workflows inside Gmail. PO confirmations, vendor queries, compliance notifications - automated end-to-end.

Outlook Connector

Extend the same email-driven automation to Microsoft Outlook environments. One agent configuration, both email ecosystems covered.

Web Search Connector

Give any agent live web search capability. Market data, supplier checks, regulatory lookups - agents answer beyond their trained knowledge, in real time.

It's Not About Which LLM.

It's About Knowing Which Class Fits.

This is the reasoning Avagama applies before any model is ever chosen — weighing your process complexity, governance requirements, and cost constraints.

Foundation-Class Models

BEST FOR: Complex reasoning, multi-step decisions, contract interpretation, nuanced compliance

WHEN: High decision complexity · irregular exceptions · senior-judgement equivalent required

Open-Source Mid-Size Models

BEST FOR: High-volume structured decisions, invoice classification, data extraction at scale

WHEN: 50,000+ transactions/month · low latency · cost-sensitive · structured input

Specialized Fine-Tuned Models

BEST FOR: Industry-specific terminology, regulated process logic, proprietary workflows

WHEN: Regulated industries · data sovereignty · confidential training data

Small / Edge-Class Models

BEST FOR: Lightweight classification, low-latency extraction, CPU-only inference

WHEN: Latency-critical infra · CPU-only environments · narrow single-purpose tasks

This is decision intelligence, not a model catalogue — the reasoning holds even as the underlying models change.

17+ Enterprise Capabilities. Day One. No Configuration Fees.

Built for finance, compliance, and audit - not just engineering teams.

Guided Customer Onboarding

Workspace configuration, user provisioning, and first agent deployment — structured and fast.

Multi-Tenant Architecture

Full data, user, agent, and configuration isolation within a shared, cost-efficient platform.

Role-Based Access Control (RBAC)

Granular, role-driven permissions — critical for security and compliance sign-off in regulated environments.

AI Observability & Tracing

Full execution tracing for every agent interaction — debug, audit, and investigate at any point.

Structured Evaluation (LLM-as-Judge)

Automated output quality assurance that scales with usage — defensible evidence for audit teams.

Token & Cost Dashboard

Per-user, per-module cost analytics — the financial transparency CFOs require before scaling.

Audit Trails on Every Execution

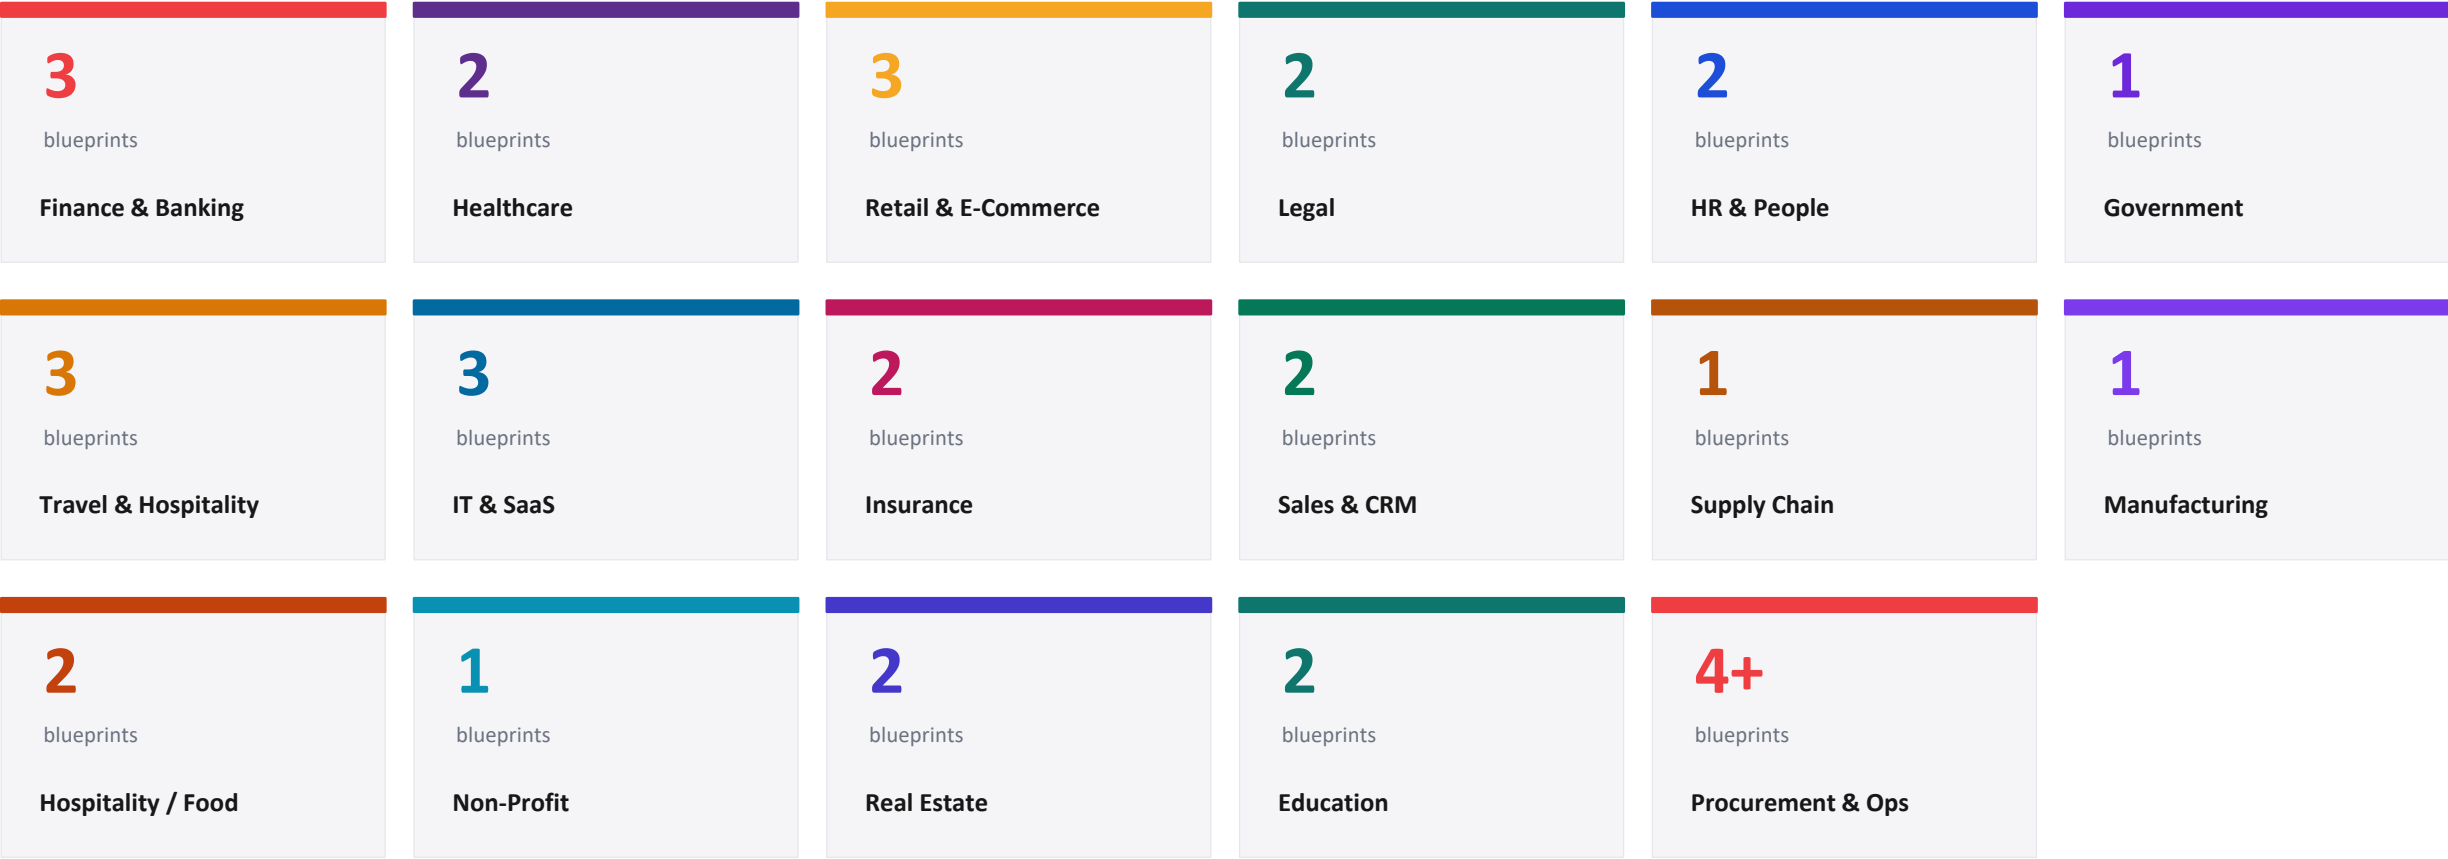
Every agent action is logged, traceable, and reportable — governance built in, not bolted on.

Confidence Thresholds & HITL

Define when agents escalate to humans. Never deploy autonomous execution without the right controls.

Running in Days. Not Months.

Every blueprint ships with a pre-defined role, goal, and LLM recommendation — not a blank canvas. Most customers have their first agent running within hours of onboarding.



Blueprint library is expanding every quarter - driven by live enterprise deployments. The more we deploy, the smarter the recommendations get.

Real Agents. Real Outcomes — Procurement & Finance.

PO & Contract Match Agent

Continuously matches purchase orders against contracts and flags discrepancies before payment — running 24/7.

OUTCOME:

90%+ reduction in manual PO reconciliation. Discrepancies caught at source — before vendor disputes, credit notes or margin leakage.

PO-to-Sales-Order Agent

Extracts structured data from incoming POs and auto-generates the matching sales order — eliminating manual re-keying.

OUTCOME:

Order processing from hours to minutes. Transcription errors eliminated. Sales ops redirected from data entry to customer engagement.

GRN Automation Agent

Generates Goods Receipt Notes from shipment data the moment stock arrives — closing the gap between physical receipt and system record.

OUTCOME:

Inventory reconciliation completed same-day. Stock discrepancies reduced significantly. Finance closes AP faster with accurate GRN data.

Master Data Management Agent

Cleanses, deduplicates and standardises Material and Service Master data at enterprise scale — continuously, not just at project milestones.

OUTCOME:

Eliminates silent errors from bad master data across procurement, finance and reporting. One source of truth, maintained automatically.

Real Agents. Real Outcomes — Vendor, Sales & Compliance.

Vendor Query Bot

Handles vendor queries directly - pricing, order status, invoice questions — without routing through procurement staff.

WHY IT MATTERS:

Vendor response time from 2–3 days to under 5 minutes. Procurement reclaims 30–40% of time on repetitive query handling. Now expanding to buyer-facing inbound supplier onboarding.

Pricing Agent

Delivers instant, defensible pricing responses to any customer query - consistent, policy-compliant quotes without waiting for approval chains.

WHY IT MATTERS:

Pricing delays - a top reason deals stall - eliminated. Sales teams close faster. Average quote turnaround reduced from 24+ hours to under 2 minutes.

TDS Compliance Agent

Production-grade compliance agent managing TDS workflows under the Income Tax Act, 1961 - deduction calculation, certificate generation and filing support, automated.

WHY IT MATTERS:

Manual TDS compliance effort reduced significantly. Regulatory risk lowered through consistent, auditable processing. Already extending to Income Tax Act 2025.

What's Coming — Driven by Customer Conversations.

The direction: deeper discovery, faster time-to-value, and no-code agent building for business teams — not just developers.

Now — Live

- 31 Blueprints across 17 industries
- Multi-framework Agent Builder
- LLM / SLM Fitment Scoring
- Execute Agent functionality
- Enterprise governance & RBAC
- Slack, Gmail, Outlook connectors
- Full observability & tracing
- Token & cost dashboard

Sep 2026

- Unified Dashboard (all agents & KPIs)
- No-Code Agent Builder
- AI-Assisted RFP Creation
- Ambition Matrix — AI readiness scoring
- Tech Stack Discovery module
- Expanding Blueprint Library (40+)
- Agent Auto-Build from Blueprints

Dec 2026+

- Process Mining Integration
- Microsoft-Native Agent Creation
- Advanced multi-agent orchestration
- Predictive ROI modelling
- Industry-specific LLM fine-tuning
- Partner agent marketplace

Done Experimenting. Ready to Execute.

Globally Unique Intelligence

The only platform that recommends the right LLM for each specific process — objectively, automatically, auditably. No competitor offers this.

Discovery to Execution — One Platform

From which process to which model to which agent running in production. No stitching together UiPath, IBM, and Celonis. One decision. One platform.

31 Blueprints — Live in Hours

Not templates. Pre-validated, industry-specific intelligence. Your first production agent in hours, not months. The library grows with every enterprise we deploy.

Enterprise-Grade, Not a Demo

Production agents running on real infrastructure with governance, observability, audit trails, and cost transparency built in from day one.

Total Model Freedom

OpenAI, Anthropic, Mistral, NVIDIA, your own models — choose based on governance, cost, and compliance. Switch as the market evolves. Your agents stay yours.

Measurable Business Outcomes

Every use case enters with a defined ROI target. Every agent runs with cost and quality visibility. Finance teams can account for every rupee of AI investment.

Avagama is for enterprises that are done experimenting and ready to execute — structured AI investment, governed for scale, measured by real business outcomes.

Let's Talk.

Start the conversation about where AI creates real value in your enterprise.

EMAIL: info@avaali.com

AVAGAMA: www.avagama.ai

AVAALI: www.avaali.com

VELOCIOUS: www.velocious.ai

FOLLOW [LinkedIn](#) • [X / Twitter](#) • [YouTube](#)

Avagama in 30 Seconds

- The world's first discovery-to-execution enterprise AI lifecycle platform
- LLM / SLM recommendation per process - no competitor offers this
- 31 blueprints across 17 industries - live in hours, not months
- Bring your own LLM - or let Avagama recommend the optimal one
- Production agents running with full governance and audit trails
- Proven in procurement, finance, compliance and operations