

ZenseAI. AgentMesh

Enterprise AI platform for
context-engineered agentic AI

 Brochure

Agentic AI, from pilot to production.

Most enterprises have AI pilots. Few have production-grade agents. ZenseAI.AgentMesh is the enterprise platform that closes the gap — with the integrations, governance, and quality engineering enterprises actually need to run agents at scale.

40+

Ready-to-use agents

**6–8
weeks**


Poc to production

30%

Cost reduction

50%

Productivity gain



01 — The Problem

Stuck in the PoC trap

Pilots succeed in isolation but stall on the path to production — blocked by integration gaps, governance risk, and quality unknowns. The result: little real adoption, automation, or business impact.

- Fragmented data ecosystems that prevent contextual awareness
- Absence of enterprise-grade guardrails and auditability
- Inconsistent performance across standalone, single-task chatbots
- Inability to securely execute multi-step actions in core systems

02 — The Platform

One production-ready platform

ZenseAI.AgentMesh is an AI-first platform to discover, build, deploy, and govern autonomous agents at scale — engineered for the last, hardest mile of agentic AI.

- **Discover & build fast** — Low-code, no-code studio with a library of 40+ vertical and horizontal agents.
- **Connect cleanly** — Native MCP / A2A connectors to your core systems, data, and other agents.
- **Govern with confidence** — 5-layer observability, financial guardrails, and regulatory controls built in.

From idea to integrated. From integrated to industrialized. AgentMesh's five engineered pillars turn experimental agents into enterprise-ready ones — without the rebuild.

03 — Outcomes you can achieve

30%

Lower cost by automating complex processes through AI orchestration

50%


Productivity improvement across orchestrated multi-agent workflows

Open

Architecture that avoids technology lock-in and future-proofs AI investment

5-layer

Observability with built-in financial & regulatory compliance controls



Five pillars, engineered for production.

P1 · Build



Low-code / no-code studio

Build agents in hours, not quarters. 50+ agents across BFSI, Retail, Manufacturing, Healthcare, Tech & Media — from KYC and fraud detection to predictive maintenance.

P2 · Reuse



100+ plug-and-play components

Build agents in hours, not quarters. 50+ agents across BFaA growing library of modular, pre-built agents for compliance, risk analytics, data orchestration, and reconciliation — reusable across every department and vertical. SI, Retail, Manufacturing, Healthcare, Tech & Media — from KYC and fraud detection to predictive maintenance.

P3 · Connect



MCP / A2A & context engineering

Native Model Context Protocol and agent-to-agent support means agents connect cleanly to your systems, tools, and other agents — with the right knowledge and guardrails.

P4 · Trust



Observability & trust

Traceable explainability for every decision, real-time security guardrails, transparent token-cost tracking, and audit trails with automated drift detection.

P5 · Quality



Quality engineering for AI

Automated test cases validate every agent workflow and action continuously — functional, behavioral, regression, and edge-case testing built into AgentMesh, not bolted on. Because production AI demands production discipline.

Our proven journey to Industrialized AI



1.

AI Maturity Check & Pilot Gap Assessment

Understand your readiness and prioritize impact.



2.

Deploy Modular AI Agents

Pre-built agents configured for your workflows.



3.

Integrate Systems & Data

Seamless integration with your ecosystem.



4.

Real-time Observability

End-to-end visibility and insights.



5.

Test & Ensure Quality

Automated testing and quality guardrails.



6.

Optimize & Scale

Continuous improvement and enterprise-wide scale.

6-8 WEEKS

Faster to value. Built for trust. Ready for scale.

>60%

Straight-through processing; document ingestion and risk scoring run concurrently.

KYC · Global
Retail Bank

50%

Loss reduction — adaptive graph analysis and ML scoring replaced static rules.

Fraud ·
Global Insurer

55%

Auto-approval; claims processed in 3 days vs. 2 weeks with AI workflows.

Claims ·
UK Insurer

60%

Fewer manual efforts via Gen AI-powered intelligent document processing.

IDP · Financial
Advisor

60s

Per compliance analysis with multi-agent regulatory orchestration.

Compliance ·
Sa Bank

06 — Why Zensar

True
interoperability

Open ecosystem

Native MCP connectors for SAP, Salesforce, ServiceNow, Snowflake, and Databricks — plus on-prem LLM support. Zero vendor lock-in.

Claims ·
Uk insurerregulatory
defensibility

Built for examiners

Cryptographically chained WORM audit logs and per-decision SHAP explanations aligned to SR 11-7, the EU AI Act, and RBI FREE-AI.

Flexible
deployment

Any cloud, any core

Azure-native with Azure OpenAI, or bring your own cloud. Technology-agnostic and open-source friendly by design.


Ready to deploy Agentic AI at scale?

Let's build together.

AgentMesh gives you the agents, orchestration, governance, and observability to go from PoC to production in weeks — not quarters.

Start your first use case →

www.zensar.com





zensar
An  **RPG** Company

At Zensar, we're 'experience-led everything.' We are committed to conceptualizing, designing, engineering, marketing, and managing digital solutions and experiences for over 145 leading enterprises. Using our 3Es of experience, engineering, and engagement, we harness the power of technology, creativity, and insight to deliver impact.

Part of the \$4.8 billion RPG Group, we are headquartered in Pune, India. Our 10,000+ employees work across 30+ locations worldwide, including Milpitas, Seattle, Princeton, Cape Town, London, Zurich, Singapore, and Mexico City.

For more information, please contact: info@zensar.com | www.zensar.com