

Policy Meets the Proprietor

A Decentralized Agentic
Architecture for India's 110 Million
MSMEs



Author:
AMODINI A P
AI Research Fellow (May-Dec 2025)

People+ai,
EkStep Foundation

Introduction

India's MSME sector employs over 110 million people and contributes roughly **30% of the country's GDP**. The government has responded with an extensive architecture of support: credit guarantee schemes, export incentives, technology upgradation funds, procurement mandates, and district-level programmes that vary by state and sector.

The Udyam portal alone has registered tens of millions of enterprises. By most measures, the policy intent is serious and the financial commitments are substantial.

And yet, a NITI Aayog report on MSME competitiveness found that the effectiveness of these schemes is consistently hampered by low awareness among the very businesses they are designed to serve. Research published in 2023 found that despite the availability of welfare-related initiatives, many small entrepreneurs lacked the institutional knowledge or resources to access them.

A 2024 study identified market access constraints and low technology adoption as persistent barriers, noting that many enterprises remain either unaware of state government schemes or unable to navigate them without external help.²

The gap, in other words, is not between policy and funding. It is between policy and comprehension. A grocery store owner in Uttar Pradesh and an electronics manufacturer in Pune face the same structural problem. The information they need exists, but it is fragmented across ministries, portals, and gazettes, rarely contextualized to their specific geography, sector, or size, and almost never available in the language they work in.

A decentralized agentic architecture

The instinct when faced with fragmented data is to centralize it: build one portal, one database, one interface. This approach runs into predictable problems. Governance conflicts arise when data custodians are unwilling to surrender ownership.

Regulatory nuance gets flattened when a scheme valid in Maharashtra is treated the same as one valid in Tamil Nadu. And large centralized systems tend to be slow to update, creating exactly the kind of temporal mismatch that makes AI advice dangerous in policy-sensitive domains.

The prototype we developed at People+ai, which we called a Decentralized Agentic Network, takes a different architectural approach. Autonomous Regional Agents operate within specific domains or geographies, managing local data interfaces and respecting governance boundaries.

When a query requires cross-domain intelligence, agents coordinate with each other through structured protocols rather than through a central repository. Data remains with its rightful custodians; intelligence travels.

To enable this coordination, we adopted the Model Context Protocol (MCP) as the connective tissue between components. An access portal connects to MCP clients, which interface with domain-specific MCP servers registered within a shared registry.

When a user query is raised, the relevant agent dynamically retrieves contextual information from appropriate domain servers and returns actionable intelligence rather than raw datasets. The architecture is modular, discoverable, and designed to scale without tightly coupling systems.

From data to decision

The deeper design principle behind this work is that the goal of an intelligence layer is to serve a decision, not to display information. A business owner asking whether they qualify for an export incentive under a specific scheme does not need a link to the gazette notification. They need a yes or no, with the conditions attached and the application pathway explained.

Building toward that outcome requires more than aggregating data. It requires agents that are grounded in verified sources, stable under user challenge, and sensitive to the temporal and jurisdictional specifics that determine whether a scheme applies.

These are the same safety properties that People+ai and Anthropic have been developing together in other applied contexts, and they are directly relevant here.

The MSME context is, in this sense, a productive stress test. The information environment is genuinely complex, the stakes of inaccurate advice are real, and the user base spans levels of digital literacy, languages, and institutional contexts that no single model trained on English-language data can handle without deliberate design.

Acknowledgements

This blog is part of People+ai's Data Series, documenting applied research on Already data infrastructure for population-scale impact.

