

SoftServe

Agents4Everything

One platform. Every agent. No rebuilds.

Most enterprises don't have an AI problem. They have a fragmentation problem — a different tool for every agent, a different interface for every team, and a rebuild cost every time someone asks for something new. Agents4Everything is built to end that.

THE PROBLEM AI DEPLOYMENTS KEEP RUNNING INTO

AI adoption stalls not because of a shortage of capable models, but because of the infrastructure required to deploy and coordinate them. Each new agent means a new interface. Each new use case means a new integration. By the time a pilot reaches production, the team has rebuilt the same scaffolding three times.



LENGTHY TIME-TO-MARKET

Manual data preparation and model training cycles consume much of a project's timeline before a single user sees a result.



INTEGRATION COST ON EVERY ITERATION

Combining different Gen AI tools and isolated agents under one interface requires bespoke development that rarely transfers to the next use case.



UX REBUILT FOR EVERY DEPLOYMENT

UI development and customization for each new client use case consumes a disproportionate share of delivery budgets.

NO FEEDBACK LOOP ACROSS AGENTS

WHAT AGENTS4EVERYTHING IS BUILT TO DO

Agents4Everything is a unified, polymorphic AI platform that lets enterprises rapidly assemble, orchestrate, and scale AI agents across business workflows without rebuilding interfaces or integrations for every new use case. It combines a plug-and-play agent architecture with a fully customizable widget-based UI, so every agent runs through a single adaptive interface that reshapes itself around the work.

BUSINESS IMPACT

30%+

faster pilot delivery cycles
using prebuilt agents and
UI components

50%+


faster production rollout
with production-ready
deployment out of the box

* Performance estimates are based on benchmark data and SoftServe pilot projects.




HOW IT WORKS

AGENTS CORE



A suite of prebuilt agents with an orchestration layer that coordinates every capability. Multimodal RAG handles documents, images, tables, graphs, and video. Talk2DB translates natural-language questions into SQL and generates charts on demand. The ML Agent runs analytics and machine learning pipelines in direct response to a prompt. New agents in development will plug in via MCP protocol without touching what already works.

POLYMORPHIC UI




A widget-based workspace that assembles only the components each use case requires. Layouts are configured manually or from a single natural-language prompt. Prebuilt views exist for different user personas. Conversation history and workspace state persist across sessions. The interface reshapes itself around the work. Every time.

PLATFORM CAPABILITIES

CAPABILITY	WHAT IT DOES
SINGLE ENTRY POINT	All agents (Multimodal RAG, Talk2DB, ML Agent, and any custom agent) are accessed through one unified interface. No context switching. No separate tools.
PROMPT-DRIVEN UI UPDATES	A natural-language prompt creates or reconfigures a multi-widget dashboard. The UI Facilitator reads the query, the current layout, and session history, then decides which agents to invoke and where to place their outputs.
WIDGETS-BASED WORKSPACES	Grid-anchored, drag-and-drop widgets with snap-to-grid behavior and resize controls. Users add, rearrange, and remove widgets manually or via prompt. Preferred layouts are saved and persist across sessions.
META AGENT ORCHESTRATION	The AI Orchestrator routes queries across agents using FastMCP and LangGraph. The UI Facilitator acts as a bridge, translating orchestrator outputs into workspace changes, adding new widgets, removing stale ones, and populating data without manual configuration.
PLUG-AND-PLAY ARCHITECTURE	Every capability is a composable agent. New agents extend the platform via MCP protocol. Nothing existing breaks. A meta-agent coordinates all interactions.
SESSION CONTINUITY	Conversation history and workspace state persist across sessions. Users return to exactly where they left off, with full context intact.
AI-DRIVEN OR MANUAL LAYOUT	Layouts can be built and modified by prompt or by hand. Both paths produce the same result, a workspace configured precisely for the task at hand.
PRODUCTION-READY INFRASTRUCTURE	Runs on AWS. Frontend: Next.js + React with SSR. Backend: FastAPI BFF. Identity: Keycloak OAuth2/OIDC. Orchestration: LangGraph. Vector search: OpenSearch. Observability: Langfuse. Analytics: ClickHouse. Caching: Redis. LLM providers: OpenAI and Anthropic



SELECTED USE CASES




DATA EXPLORATION VIA CHAT - A business analyst types a question about Q3 performance. Talk2DB generates the SQL, queries the database, and populates a table and chart widget in the same workspace—in under a minute. The data team never gets the ticket.



DOCUMENT INTELLIGENCE - A user asks the Multimodal RAG agent to extract key findings from documents. The agent processes text, tables, images, and graphs across every file simultaneously. Structured results surface into a report widget without opening a single document.



ML-DRIVEN FORECASTING - A marketing specialist types a forecasting request in plain language. The ML Agent selects the appropriate pipeline, runs it, and renders an interactive chart widget with historical data overlaid. No code written. No pipeline configured.



MULTI-AGENT DASHBOARD ASSEMBLY - An entire dashboard is assembled using prompts and combining RAG-sourced document summaries, database-queried metrics, and ML-generated forecasts across separate widgets, automatically arranged, populated, and ready to share.

WHY SOFTSERVE



Product-grade platform, not a one-off. Agents4Everything is built as a repeatable, scalable product. Every client deployment reuses the same core and extends it rather than rebuilding from zero.



End-to-end delivery. AI, ML, UX, backend, and enterprise integration expertise under one team. Accelerator-based delivery compresses implementation time and reduces integration effort across projects.



Built-in Gen AI tooling. Insight Engine, Talk2DB, and ML Agent are prebuilt and production tested. MCP orchestration templates follow best-practice patterns from the ground up.



Flexible engagement. Advisory, model build, validation, and complete end-to-end delivery. Fixed-price custom agent development. Engagement models aligned to where the client is, not where a contract assumes they should be.

TECHNOLOGY AND DEPLOYMENT

Agents4Everything runs on Amazon Web Services with full support for Microsoft Azure configurations. The platform can be easily integrated with enterprise identity providers via Azure AD and Okta, connects to OpenAI and Anthropic LLM providers, and supports enterprise data storage including Snowflake and SQL databases.

Deployment options cover cloud, on-premises, and hybrid configurations. The modular architecture supports extension to new workflows without rebuilding existing integrations.

Focus industries: cross-industry, with production deployments in manufacturing, retail, financial services, healthcare, and energy.

ABOUT SOFTSERVE

SoftServe is a premier IT consulting and digital services provider with more than 30 years of experience delivering superior digital solutions to enterprise industries including high tech, financial services, healthcare, life sciences, retail, energy, and manufacturing. SoftServe architects and executes mature and innovative capabilities across digital engineering, data and analytics, cloud, AI/ML, and enterprise integration.

Contact SoftServe to run a live Agents4Everything session against a data environment that looks like yours. You leave with a clear pilot scope, an integration map, and a fixed-price implementation path from platform to production.

GET STARTED →