

Agent Studio

The control plane for enterprise AI agents

Create, configure, test, and orchestrate AI agents across enterprise data, tools, and systems, with the governance and observability required to run agents safely in production.

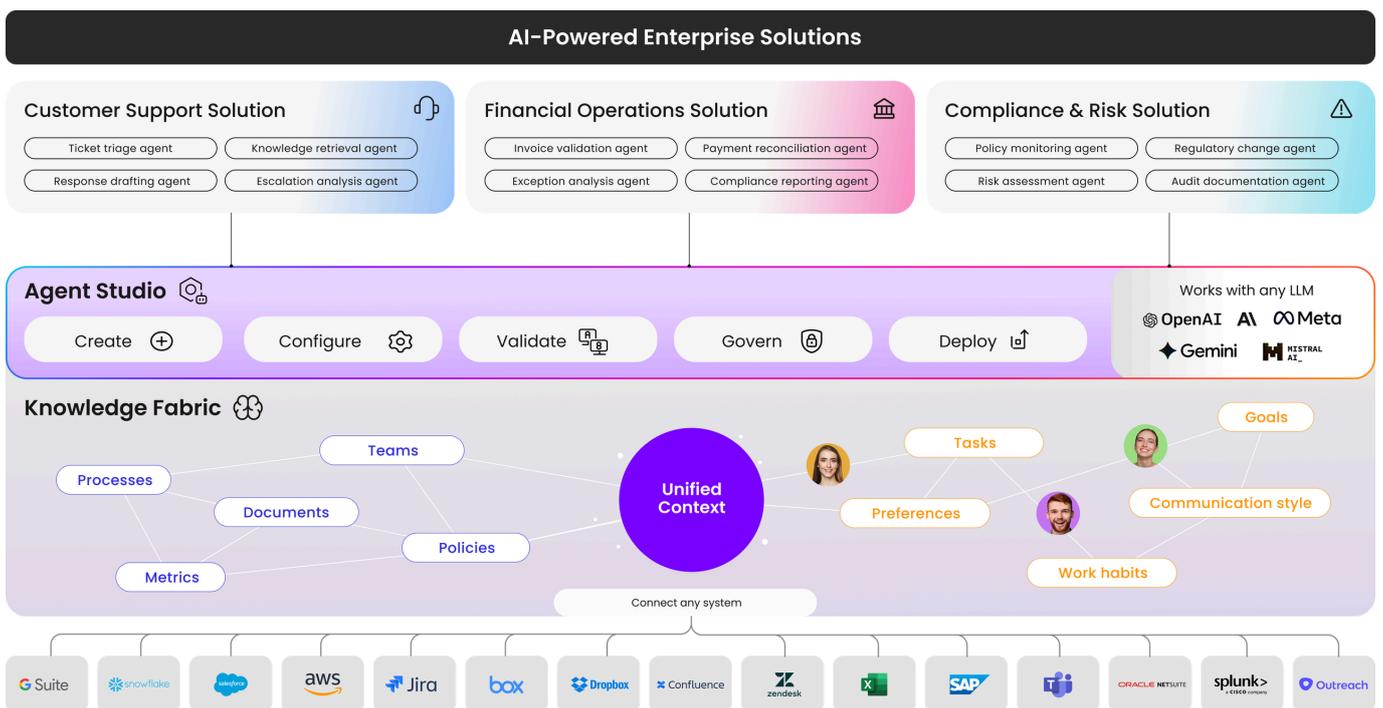
Operating AI agents across the enterprise is complex

Prototyping an AI agent is easy. Running agents across enterprise systems is not. Enterprise agents must interact with internal tools, operate on sensitive data, follow governance policies, and remain observable and auditable in production environments.

Organizations rarely deploy a single agent. Instead, they introduce dozens across business functions, each responsible for tasks like research, approvals, reporting, analysis, or customer interactions. As agents automate workflows across enterprise systems, organizations need a centralized control plane to design, operate, and govern them at scale.

Create, configure, and operate AI agents with Unframe Agent Studio

Enterprise solutions are increasingly powered by dozens of AI agents, each responsible for a specific workflow. Agent Studio is the control plane where those agents are created, tested, and governed. Teams use it to configure how agents operate, validate their behavior before deployment, and observe how they run across enterprise systems, making it possible to operate agents safely at scale.



How Agent Studio Works



Create and configure agents

Agent Studio enables teams to rapidly create AI agents using predefined templates aligned to enterprise workflows. Organizations can control which tools agents can use, what knowledge and data they can access, which models power their interactions, and how they interact with enterprise systems.



Extend agents with enterprise tools

Enterprise agents must interact with real systems to deliver value. Unframe enables organizations to extend agent capabilities through MCP-compatible servers, allowing teams to connect internal APIs, external services, and custom enterprise functions.



Ground agents on enterprise context

Agents run on your instance of the Unframe Knowledge Fabric, a shared context layer connecting your enterprise data, systems, and organizational knowledge. Teams can also shape this context directly in natural language – defining terms, adding guidance, and teaching the system how information should be interpreted. Agents automatically inherit this knowledge as they operate across enterprise systems.



Test agents before production

Before deployment, agents can be tested interactively to validate their behavior and outputs. Teams can simulate interactions, evaluate how agents query data, refine agent configuration, ensuring agents behave reliably before they are deployed into production workflows.



Observe and govern agent behaviour

Unframe provides deep observability into agent execution, giving teams clear visibility into how agents operate across enterprise systems. Administrators can inspect which agents ran, the steps they took, the queries they executed, and the tools and systems they accessed. Policies for data access, system permissions, and operational guardrails are defined upfront, ensuring agents operate within enterprise governance standards.



Deploy agents across the business within hours

Once validated, agents can be deployed across enterprise environments through APIs and integrations. Because agents operate on the shared Knowledge Fabric and platform infrastructure, new agents can be introduced quickly without rebuilding integrations or infrastructure.

Configure model providers centrally

Define which models power agent interactions and which providers are available across the environment. This ensures agent behavior aligns with internal standards for cost, performance, quality, and governance. Agent Studio's multi-model architecture standardizes model access across the organization while keeping control centralized.

Personalize agent experiences for every user

Unframe supports user-level memory and personalization. Teams can define user-specific preferences – such as preferred language or response style – through direct interaction. Agents adapt to individual users while keeping those adjustments transparent and governed.