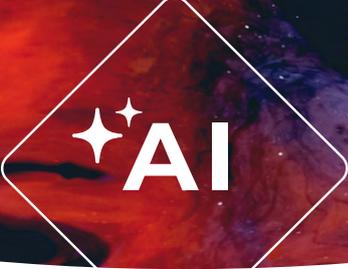
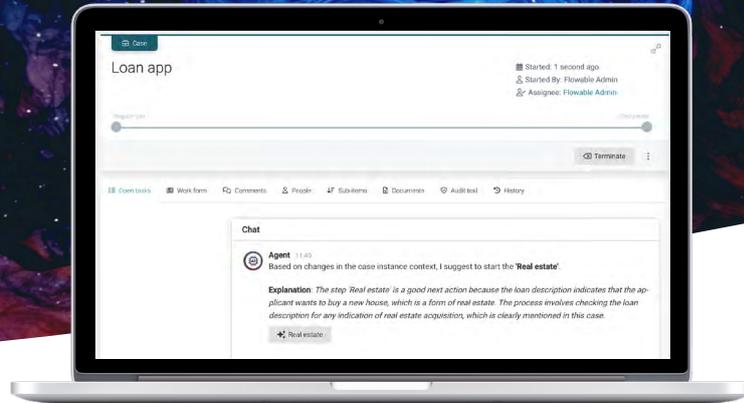


Flowable AI Studio



AI



Flowable AI Studio provides a comprehensive suite of capabilities to embed intelligent automation into business workflows. It enables the design, orchestration, and execution of AI-powered logic across processes, cases, and user interactions.

Agentic Case Platform

The Agentic Case Platform is Flowable's flagship innovation for intelligent case management. It brings together AI agents, human users, and business processes into a unified, orchestrated environment that operates autonomously and adaptively. Designed to handle complex workflows with minimal oversight, the platform functions 24/7, anticipating issues and resolving them proactively based on historical data, LLM training, and real-time context.

The platform is agnostic to the LLM and AI vendor choices made to provide a foundation for any multi-agent orchestration requirements. AI agents can be modeled as part of the case or process model and therefore directly benefit from the rich modeling palette of CMMN and BPMN. But also, as part of a work or task form AI agents can be easily integrated and have the ability to present dynamic suggestions or do instant summaries or translations.

Flowable provides full flexibility in the level of autonomous AI logic that is the best fit for the use case. Dynamic AI prompts can be evaluated against a LLM or external AI service of choice and the prompt outcome can be used in the further case or process execution. It is also possible to use the AI outcome as a suggestion and let a knowledge worker evaluate the suggestion and take the final decision.

Core Capabilities

- **Autonomous Execution**
AI agents can independently manage tasks, reduce manual effort and improve accuracy.
- **Human-in/on/out-the-loop**
The platform supports flexible collaboration modes, allowing humans to intervene, supervise, or delegate entirely to agents.
- **Scalability**
It handles large volumes of cases simultaneously, enabling organizations to scale operations without increasing overhead.
- **Governance & Compliance**
Every agent action is logged, ensuring full auditability and transparency for regulatory needs.
- **Guardrailing**
The data, tooling and execution context an agent has access to can be fully controlled and by inheriting a user's permission set when interacting with an agent prevents data leakage or misuse of agentic power

Flowable AI Studio

Agent Types

- Orchestrator Agents**
 Drive workflows by coordinating other agents, sub-processes, services and human tasks. They interpret case context and determine next steps autonomously or suggestively.
- Utility Agents**
 Perform specialized tasks such as data execution or transformation and enrichment or validation as well as content generation. These agents are lightweight and highly reusable.
- Document Agents**
 Classify, OCR, extract, summarize, and translate content from unstructured documents. They are central to IDP (Intelligent Document Processing) workflows.
- External Agents**
 Integrate with third-party AI services (e.g., Salesforce, AWS Bedrock, Azure AI, Claude) to extend capabilities beyond the Flowable ecosystem.

Each agent type is designed to interact with cases, processes, and external systems, enabling dynamic orchestration and intelligent decision-making across the enterprise

AI-Powered Process & Case Modeling

Flowable AI Studio simplifies and accelerates workflow creation by allowing users to generate processes, cases, and forms from natural language prompts. This feature helps overcome the “blank canvas” problem and supports rapid prototyping. AI-generated models are fully customizable and can incorporate existing service definitions and data structures, ensuring continuity and alignment with business goals.

Multi-Agent Orchestration

Multi-agent orchestration allows multiple AI agents

to collaborate within a single workflow. Agents can pass data, invoke functions, and coordinate actions based on BPMN and CMMN standards. This enables both deterministic and adaptive execution paths. The orchestration engine supports internal and external agents, ensuring flexibility and extensibility across enterprise systems.

Human-AI Collaboration (Autonomous Aspect)

Flowable supports autonomous agent execution while maintaining human oversight. Agents can operate independently or suggest actions for human approval. This hybrid model ensures that critical decisions remain under human control while routine tasks are automated. The platform’s design encourages trust, transparency, and ethical AI usage.

MCP (Model Context Protocol)

MCP is the foundational protocol that enables structured communication between AI agents and traditional components within Flowable workflows. It ensures that agents operate with full awareness of process state, user context, and business rules. MCP supports real-time orchestration, allowing agents to reason across data sources and maintain continuity across tasks. It’s embedded deeply into the platform’s architecture, making agentic automation predictable, scalable, and compliant.

Retrieval-Augmented Generation (RAG)

Flowable AI Studio supports Retrieval-Augmented Generation (RAG), allowing AI agents to incorporate enterprise knowledge – such as documents, emails, and structured data – into their responses. This enables more accurate and context-rich interactions, aligning outputs with internal guidelines, handbooks, and regulations.

Instead of relying solely on foundational models, RAG empowers agents to retrieve relevant content from your organization’s knowledge base. This improves compliance, reduces hallucinations, and ensures that AI-driven decisions are grounded in real business data.

Flowable AI Studio

AI-Enhanced Forms

Flowable allows AI to be embedded directly into forms, enabling conversational interfaces, auto-fill capabilities, and dynamic field suggestions. This enhances user experience and reduces friction in data entry. AI-enhanced forms can also trigger workflows based on user input, making them a powerful tool for front-line automation.

AI Service Registry

The AI Service Registry provides an open and flexible architecture for integrating various LLMs and AI services. Users can register services like Azure AI, AWS Bedrock, Claude, and OpenAI, and invoke them via REST or scripting. This modular approach supports multi-LLM switching based on accuracy, speed, or cost requirements.

AI Agent Audit Trail

All interactions with AI agents are logged in the Agent Exchange Table, providing a complete audit trail of messages sent and received. This ensures transparency, supports compliance audits, and helps explain AI-driven decisions. The audit trail includes prompts, responses, context, and execution metadata, offering full traceability across workflows.

AI-Driven Document Insights

Flowable's AI capabilities streamline document handling by offering automatic summarization, sentiment analysis, and conversational querying. These features help users make smarter decisions and reduce time spent navigating complex content.

Documents can be translated automatically, simplifying multilingual collaboration and improving accessibility across regions. AI agents also classify documents and extract structured data—such as names, dates, and amounts—making unstructured content actionable within workflows.

These capabilities are strengthening Flowable's Intelligent Document Processing (IDP) functionality, enabling faster, more accurate automation in areas like onboarding, compliance, and customer support.

Generative AI Modeling

Generative AI Modeling in Flowable AI Studio is designed to accelerate the creation of intelligent workflows, applications, and automation logic using natural language prompts. It addresses the "blank canvas" problem by enabling users to start with AI-generated models that are fully editable and context-aware. This capability is not limited to isolated tasks – it spans across case models, process flows, forms, and even full applications.

Dual Modeling Modes

Flowable AI Studio supports two distinct modes of generative modeling:

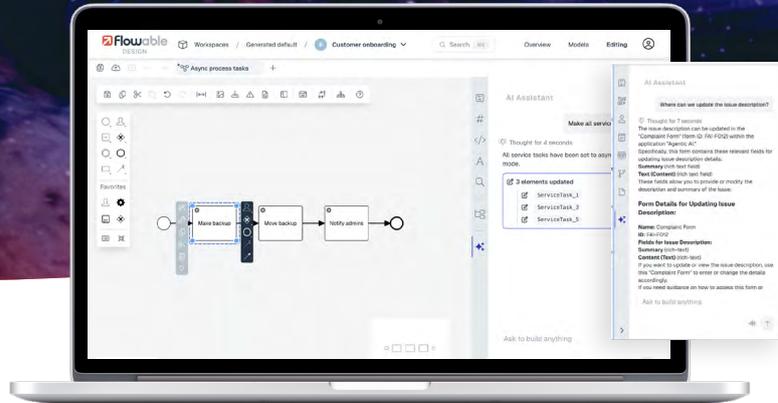
1. Full Application Generation

Users can prompt the system to generate an entire application, including case models, process flows, form structures, and data dictionaries. This is ideal for rapid prototyping or when launching new business solutions from scratch. The generated application is coherent, contextually aligned, and ready for refinement.

2. Single Model Generation

For more targeted use cases, users can generate individual models – such as a case or process – based on specific prompts. This allows for modular development and iterative refinement, where users can optimize the model through follow-up prompts and feedback loops.

Flowable AI Studio



Iterative Refinement

Unlike static generation tools, Flowable AI Studio supports prompt chaining and iterative optimization. Users can refine the output by providing additional context or corrections, allowing the AI to regenerate improved versions of the model. This feedback-driven loop ensures that the final model aligns precisely with business requirements and user expectations.

Integration with Existing Assets

Generated models can incorporate existing service definitions, data structures, and business rules. This ensures continuity and leverages prior investments in modeling. It also allows users to build on top of what already exists, rather than starting from scratch every time.

Governance and Control

All generative outputs are subject to Flowable's governance framework. This includes:

- Retrieval-Augmented Generation (RAG) to ensure AI uses approved and accurate information sources.
- Audit Trails for every AI-generated model, capturing prompt history, decision logic, and user interventions.
- Agent Autonomy Settings to define how much control AI agents have over model creation and execution.

Generative AI Modeling transforms how enterprises approach workflow automation. It reduces time-to-value, democratizes access to modeling capabilities, and ensures that even non-technical users can contribute meaningfully to solution design. By embedding AI into the modeling lifecycle, Flowable empowers teams to innovate faster while maintaining full control over compliance, security, and business alignment.

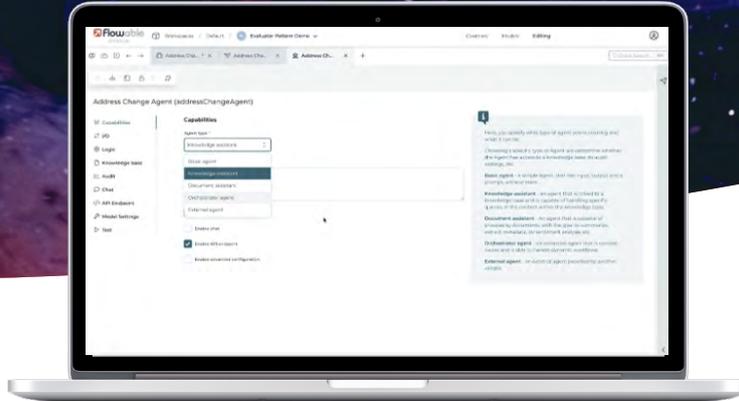
AI Governance & Guardrails

Governance features include role-based access, lifecycle management, and policy enforcement. Guardrails ensure that agents operate within defined boundaries, preventing unauthorized actions and maintaining compliance.

Secure External AI Agent Integration

Flowable supports integration with external agents like Salesforce, AWS Bedrock, and Azure AI. These agents can be orchestrated within workflows, with full auditability and permission controls. This enables cross-platform automation while maintaining enterprise-grade security.

Flowable AI Studio



AI Chat-Based Interaction*

Flowable's conversational interface spans both Design and Work.

- In **Design**, AI chat is integrated into all modeling editors (process, case, decision, form, etc.), allowing designers to add elements, update properties, or perform bulk changes via prompts.
- In **Work**, users can interact with case and process data through natural language queries, asking for summaries, next-step suggestions, or detailed information.

This conversational approach improves navigation and decision-making across the platform.

Adaptive Workflows

AI can be embedded throughout workflows using familiar patterns such as the service registry. This includes integration into forms, processes, cases, and actions – enhancing decision-making and automating tasks across use cases. Additionally, AI agents can adapt workflows in real time based on context, user behavior, and external signals. This ensures that processes remain relevant and responsive, even in dynamic environments, improving agility and reducing the need for manual intervention.

AI in Email Classification

AI agents parse and classify email content, detect intents and extract structured data for use in workflows. This supports use cases like customer onboarding, support ticket triage, and compliance monitoring. Email becomes a first-class citizen in case management.

AI-Driven Decision Support

AI agents assist in decision-making by analyzing context, retrieving relevant data, and suggesting actions. These decisions can be embedded in BPMN and CMMN models, ensuring traceability and consistency. Decision support improves accuracy and reduces cognitive load for users.

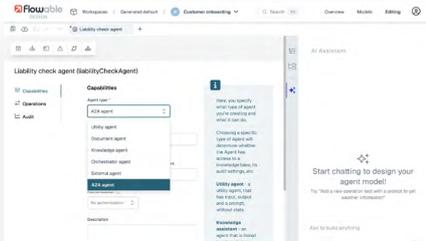
AI Studio Visual Design Tools

Flowable AI Studio includes drag-and-drop tools for modeling agents, workflows, and interactions. Users can visually define agent behavior, context, and orchestration logic. This lowers the barrier to entry and accelerates solution development.

AI-Powered App Generation

From a single prompt, users can scaffold entire applications – including models, forms, and integrations. This feature supports rapid prototyping and reduces time-to-value. Generated apps are fully editable and can be extended with custom logic.

Flowable AI Studio



External Agents with A2A Support

Flowable AI Studio supports external agents implementing the A2A specification, enabling interoperability between Flowable and third-party agent ecosystems.

MCP Support

When modeling an agent, tools can be added as additional sources of information to the LLM. These tools can include REST services or other agents, creating a flexible and powerful orchestration layer.



Agent Timeline Visualization

Flowable Control introduces a timeline view for agent instances, showing requests, responses, LLM calls, and tool invocations, with detailed request/response data for full transparency.

Agent Metrics

Agent metrics track invocations per definition and token usage, providing insights to optimize performance and costs.

AI Support in Forms

A specific AI button is available to connect with an agent model, enabling direct invocation of an agent from a form. This feature brings intelligent automation closer to end-user interactions.

Housekeeping of Agent Data

The removal of agent data is tied to regular housekeeping of case and process data, ensuring data hygiene and compliance with retention policies.

Flowable AI Studio sits on top of the [Flowable platform](#) and is licensed separately.

01/2026_Flowable AI Studio © Flowable 2026. All rights reserved. The information contained herein is of a general nature, is for informational purposes only and is subject to modification without notice. Although we strive to deliver accurate and comprehensive information, we do not give any guarantee on the information provided. The Flowable name and logo are registered trademarks.