



Flowable Open Source compared

What the open source engines provide, and what the platform adds

Flowable open source is a set of embeddable engines. The Flowable Agentic Case Platform builds on them with everything an organization needs to model, run, govern and support its work. This document compares the two, capability area by capability area.

Flowable open source and the Flowable Agentic Case Platform



Flowable open source is a set of embeddable process, case and decision engines under the Apache 2.0 license. A development team adds them to an application it builds itself and drives them from code or the REST API; there is no user interface. The Flowable Agentic Case Platform builds on those engines and goes far beyond them: more than twenty further engines, platform logic hooked into every step the engines execute, configuration tuned for high volume, the modeling environment, the end-user application, the administration console, AI agents, testing tools, and a vendor with a contractual support commitment. The open source engines contain no AI or agent capabilities at all.

One use case, or a portfolio

On a single process the surrounding layer looks like a few screens and one integration. Organizations rarely stop at one: onboarding is followed by reviews, complaints, approvals and dozens more. Permissions, audit, identity, search, retention, translations and branding then have to be solved once and behave the same everywhere, integrations have to be defined once and reused, and someone has to own all of it through every engine upgrade. The platform provides that shared layer, and each further use case costs less than the one before.

What you build, you own

Every layer written on top of the open source engines is maintained, secured and upgraded by its author, without warranty and without a vendor to escalate to. Some parts are hard to get right at all: permissions enforced on every API call, an audit trail and four-eye rule a regulator accepts, agent governance, indexing at production volume, an accessibility audit, and identity integration. The platform ships these as tested, supported product, with private correction releases and two major releases a year plus monthly patches.

Capability area	Flowable Open Source	Flowable Agentic Case Platform
Process, case and decision execution	The open source BPMN, CMMN and DMN engines, Apache 2.0	The open source engines extended with platform logic, security and data services around every step, tuned for high volume, with enterprise bugfix releases
Model authoring	Models are written by hand as XML or generated by tools	Flowable Design: 29 model types with editors, AI-assisted authoring and a Design MCP server, validation, versioning, comparison, Git, tests and controlled publishing
End-user application	Build your own on the REST API	Flowable Work: inbox, case pages, tasks, documents, collaboration, mobile; embeddable or fully headless
Forms and pages	Field definitions, nothing to render them	Form builder and renderer with validation and conditional behavior, modeled pages, user-composed dashboards, and SDKs to embed forms, cases and apps in external user interfaces

Capability area	Flowable Open Source	Flowable Agentic Case Platform
Business data	Process and case variables	Data Objects, data dictionaries with enforced constraints, master data, sequences
Integration	Written in code or configured per call with the HTTP task	Service registry: modeled, versioned, reusable integrations with authentication, mocking and OpenAPI import; more event transports
Documents	Not included	Document management with versions, metadata and previews, Word and PDF generation, document events
Search and analytics	Queries against the database	Indexed search over all work with full text, dashboards, reports and heatmaps
Permissions	Roles are recorded but not enforced; one check on API access	Security policies enforced on every API call, per instance, plus the four-eye principle
Identity	Users and groups, LDAP lookup, no single sign-on	Single sign-on with OAuth2 and OpenID Connect, LDAP and Entra ID resolved live, per-tenant configuration
AI agents	Not included	Agent engine with eight agent types, guardrails, evaluators, MCP, full tracing and cost
Administration	REST and Java APIs	Flowable Hub: instances, jobs, assisted migration of whole hierarchies, health, incidents, configuration
Testing	JUnit against the engines	Tests modeled in Design and run by Inspect, with mocking, breakpoints and assertions
Support	Community forum	Commercial support with a contractual response commitment and warranty
Releases	Public releases, roughly once a year	Two major releases a year, monthly patches, corrections before public release

Moving from open source to the platform is supported: process, case, decision and event definitions carry over unchanged, running work is preserved, and the existing REST API keeps working, now subject to permission checks and extended with the platform APIs.



flowable.com
info@flowable.com

© 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. Screenshots are illustrative and may differ from the shipped product. The Flowable name and logo are registered trademarks. Comparison v1, September 2026, for release 2026.1.