

CaseFlow

Case + Process Management

Build software for claims, applications, investigations, compliance reviews, onboarding, service requests, permits, underwriting, HR matters, or any process where a record moves through stages while accumulating people, documents, conversations, tasks, decisions, and approvals.

CASEFLOW - WALKTHROUGH SCENARIO

A compliance review enters through intake, accumulates documents and communications, hits an SLA checkpoint, gets AI blocker analysis, and advances only after the required approval.

[View evidence](#)[Start workflow](#)

1

GOVERNED PRODUCT MODEL

6

CORE CAPABILITIES

5

WALKTHROUGH STAGES

4

SEPARATED DATA / CODE PLANES

What CaseFlow is

CaseFlow packages reusable BuildWithHQ primitives into a focused case + process management product experience.

01. Product layer

A branded application experience shaped around a specific operating job.

02. Governed intelligence

AI can analyze, retrieve, recommend, and prepare actions inside the same permission envelope as the application.

03. Operational actions

Workflows, approvals, APIs, webhooks, MCP tools, and custom services move from insight to controlled execution.


04. Platform leverage

Dedicated data planes, audit evidence, exports, and optional container services stay reusable under the product.

The operational problem

CaseFlow is designed for work that becomes expensive when context, decisions, and actions are fragmented across tools and people.

Illustrative sources of operating friction



What changes

Instead of treating AI, records, workflows, permissions, and integrations as separate projects, the product keeps them inside one governed operating model.

Design goal

Reduce the distance between context, decision, and controlled action without flattening the business rules that make the workflow safe.


NOT JUST A FEATURE

A complete operating product

Turn complex cases into controlled, explainable workflows.

WHO IT SERVES

- Case worker gets a purpose-built view, authority level, and workflow.
- Reviewer gets a purpose-built view, authority level, and workflow.
- Manager gets a purpose-built view, authority level, and workflow.
- Customer gets a purpose-built view, authority level, and workflow.
- Compliance gets a purpose-built view, authority level, and workflow.
-



WHAT MAKES IT DIFFERENT

- AI is attached to governed records and evidence.
- Actions are explicit contracts, not hidden autonomous behavior.
- Permissions remain server-side screens, APIs, AI and exports.
- Code is optional - not the only match.

The six building blocks of CaseFlow

Build software for claims, applications, investigations, compliance reviews, onboarding, service requests, permits, underwriting, HR matters, or any process where a record moves through stages while accumulating people, documents, conversations, tasks, decisions, and approvals.

01 Flexible case records

Model claims, applications, reviews, investigations, requests, matters, or your own case type with the fields and relationships your industry needs.

02 Stage + SLA workflows

Define statuses, transitions, owners, service levels, timers, escalation paths, and required conditions before a case can move forward.

03 Documents + communications

Keep proposed files, generated documents, emails, messages, forms, and customer responses attached to the case and recursively related records.

04 Role-aware work queues

Route work to teams by case type, geography, customer, risk, expertise, or workload while restricting sensitive records and fields by role.

05 AI case intelligence


Summarize long histories, identify missing information, highlight blockers, prepare next actions, and give reviewers evidence-linked context before a decision.

06 Approvals + replayable audit

Require human approval for sensitive transitions and preserve who saw what, what changed, why it changed, and the evidence behind the decision.


Separated by responsibility. One product to the user.

BuildWithHQ keeps runtime records, AI context, compliance evidence, and custom execution separate while the application presents one coherent experience.




Everything important becomes connected context

The product data model is not a set of isolated modules. Records can relate recursively so users and AI can traverse the real operating context.



From input to governed outcome

A compliance review enters through intake, accumulates documents and communications, hits an SLA checkpoint, gets AI blocker analysis, and advances only after the required approval.



The five stages are illustrative - each BuildWithHQ implementation can add branches, approvals, retries, SLAs, external systems, and custom services.

1. Intake

Create the case, identify parties, collect required data, and establish initial ownership.

INPUT

What enters this stage

Identity, permitted records, related context, workflow state, and any source-specific inputs required for intake.

EVIDENCE

Evidence produced

Store the records consulted, source versions, relevant model output, workflow transition, approvals/denials, and mutation result so the operation can be explained later.

ENGINE

What BuildWithHQ does

Create the case, identify parties, collect required data, and establish initial ownership. The runtime enforces tenant, role, record, field, rate-limit, and action boundaries before the stage can progress.

EXPERIENCE

What the user sees

A focused CaseFlow screen should expose the result, why it happened, what remains unresolved, and the next safe action - without forcing the user to inspect raw infrastructure.

2. Progress

Move through controlled stages with tasks, timers, SLAs, document requirements, and escalations.

INPUT

What enters this stage

Identity, permitted records, related context, workflow state, and any source-specific inputs required for progress.

EVIDENCE

Evidence produced

Store the records consulted, source versions, relevant model output, workflow transition, approvals/denials, and mutation result so the operation can be explained later.

ENGINE

What BuildWithHQ does

Move through controlled stages with tasks, timers, SLAs, document requirements, and escalations. The runtime enforces tenant, role, record, field, rate-limit, and action boundaries before the stage can progress.

EXPERIENCE

What the user sees

A focused CaseFlow screen should expose the result, why it happened, what remains unresolved, and the next safe action - without forcing the user to inspect raw infrastructure.

3. Review

Summarize history, missing information, blockers, and evidence for the current decision.

INPUT

What enters this stage

Identity, permitted records, related context, workflow state, and any source-specific inputs required for review.

EVIDENCE

Evidence produced

Store the records consulted, source versions, relevant model output, workflow transition, approvals/denials, and mutation result so the operation can be explained later.

ENGINE

What BuildWithHQ does

Summarize history, missing information, blockers, and evidence for the current decision. The runtime enforces tenant, role, record, field, rate-limit, and action boundaries before the stage can progress.

EXPERIENCE

What the user sees

A focused CaseFlow screen should expose the result, why it happened, what remains unresolved, and the next safe action - without forcing the user to inspect raw infrastructure.

4. Approve

Require authorized human decisions for sensitive transitions and exceptions.

INPUT

What enters this stage

Identity, permitted records, related context, workflow state, and any source-specific inputs required for approve.

EVIDENCE

Evidence produced

Store the records consulted, source versions, relevant model output, workflow transition, approvals/denials, and mutation result so the operation can be explained later.

ENGINE

What BuildWithHQ does

Require authorized human decisions for sensitive transitions and exceptions. The runtime enforces tenant, role, record, field, rate-limit, and action boundaries before the stage can progress.

EXPERIENCE

What the user sees

A focused CaseFlow screen should expose the result, why it happened, what remains unresolved, and the next safe action - without forcing the user to inspect raw infrastructure.

5. Close

Capture outcome, notifications, evidence, audit history, and related follow-up work.

INPUT

What enters this stage

Identity, permitted records, related context, workflow state, and any source-specific inputs required for close.

EVIDENCE

Evidence produced

Store the records consulted, source versions, relevant model output, workflow transition, approvals/denials, and mutation result so the operation can be explained later.

ENGINE

What BuildWithHQ does

Capture outcome, notifications, evidence, audit history, and related follow-up work. The runtime enforces tenant, role, record, field, rate-limit, and action boundaries before the stage can progress.

EXPERIENCE

What the user sees

A focused CaseFlow screen should expose the result, why it happened, what remains unresolved, and the next safe action - without forcing the user to inspect raw infrastructure.

Permissions follow the user into AI, APIs and actions


One authorization model should determine what each role can read, retrieve, export, approve, or mutate.

Illustrative authority matrix - 1 restricted to 5 broad

	Read	AI context	Recommend	Approve	Execute
Case worker	5	4	2	5	1
Reviewer	5	4	4	5	4
Manager	5	4	4	2	4
Customer	5	4	4	2	1
Compliance	5	3	2	5	1

AI is a participant in the workflow - not a bypass around


The AI layer can retrieve, synthesize, rank, draft, and recommend. Server-side policies still decide what data and actions are actually available.



This pattern lets the same intelligence layer support employees, end customers, scheduled agents, MCP clients, and API-driven experiences.

Connect to the systems already around the business

API, webhooks, MCP, CLI, file import/export, and optional private services let the product fit into an existing stack instead of becoming another isolated island.



Inbound


Events, imports, messages, usage data, customer records, documents, signals, or other domain-specific inputs.

Outbound

Actions, notifications, record updates, webhooks, exports, generated files, MCP tools, or custom API calls.

Optional isolated container when the product needs


Most of the application can remain schema-, workflow-, and API-driven. When a developer needs a special SDK, proprietary algorithm, or long-running service, the isolated container provides a controlled extension point.



The central execution layer can call the declared service while the custom code remains isolated from the core platform and from other customers.


What a CaseFlow operating screen can look like

The point of the blueprint is to turn deep infrastructure into a focused daily workflow.




Every completed cycle can improve the next one

The product gets stronger when outcomes, approvals, exceptions, source quality, user corrections, and workflow performance are fed back into the operating model.



Where operating leverage can show up


The numbers below are not measured BuildWithHQ customer results. They are an illustrative model showing the kinds of friction the product is designed to reduce.



Value compounds as more work enters the governed

A staged deployment lets the team prove retrieval, routing, approvals, and actions before widening automation authority.

Illustrative percent of target workflow running through the product



Phase 1 - Assist

Read-only context, search, summaries, classification, and drafts. Human users make every consequential decision.

Phase 2 - Governed action

Add explicit tools, approval thresholds, scheduled runs, and carefully bounded auto-actions where the evidence is strong.

Six ways to package CaseFlow

One BuildWithHQ blueprint can be narrowed to a vertical, customer segment, operating team, or recurring workflow and then branded as its own SaaS product.

01 Claims processing

Use the same core CaseFlow architecture, then customize terminology, records, permissions, integrations, AI behavior, and workflows for this market.

02 Compliance review

Use the same core CaseFlow architecture, then customize terminology, records, permissions, integrations, AI behavior, and workflows for this market.

03 Permits

Use the same core CaseFlow architecture, then customize terminology, records, permissions, integrations, AI behavior, and workflows for this market.

04 Underwriting

Use the same core CaseFlow architecture, then customize terminology, records, permissions, integrations, AI behavior, and workflows for this market.

05 HR matters


Use the same core CaseFlow architecture, then customize terminology, records, permissions, integrations, AI behavior, and workflows for this market.

06 Customer onboarding

Use the same core CaseFlow architecture, then customize terminology, records, permissions, integrations, AI behavior, and workflows for this market.

A role-aware experience from morning to close

Different users can work on the same underlying records while seeing different queues, context, evidence, and action rights.



The important part

The application is not forcing every role into the same dashboard. The runtime can expose the same record with different fields, related context, actions, AI tools, and approval rights based on who is looking at it.

Ways a builder could monetize the product

These are product packaging ideas, not fixed BuildWithHQ pricing. A builder can mix software subscriptions, usage, premium modules, infrastructure, and services.

Per-case volume tier

SLA / workflow add-on

Document AI usage


Customer portal module

Industry case templates

Illustrative packaging ladder: start simple, then monetize the capabilities that scale with customer value or infrastructure consumption.

A practical 90-day build sequence

The exact timeline depends on scope, but the safest pattern is to establish data + identity first, then intelligence, then actions, then expansion.




Launch authority should expand only after the evidence trail, failure handling, permission checks, and rollback path are tested.


What to measure after launch

The product should be evaluated on operating outcomes, not just model usage. Track speed, quality, safety, adoption, and economic leverage together.

Illustrative maturity scorecard



Illustrative implementation effort mix




Make the dangerous path harder than the safe path

The goal is not to promise that AI or automation cannot fail. The goal is to constrain what failure can affect, surface evidence, and preserve a deterministic way to stop or reverse the operation.




One backend, multiple ways to use it

The same product can support a visual SaaS interface, an external frontend, internal automation, developer tooling, and agent access without creating five disconnected permission models.



Productize once. Isolate customers by design.

BuildWithHQ is intended to let a builder define the product experience while tenant, user, permission, data, AI, and optional service boundaries keep each customer context separated.



Shared platform primitives - isolated customer context - configurable product experience

BUILD THIS

Start with CaseFlow. Make it yours.

CaseFlow is a product direction, not a fixed template. Use the blueprint as a starting architecture, then shape the records, screens, AI behavior, integrations, permissions, workflows, container services, and commercial model around your market.

1. Define the job

Choose the recurring operating problem and the user roles that feel it most.

2. Map the system

Define records, relationships, evidence, permissions, workflows, and required integrations.

3. Add intelligence

Layer AI and agents into the governed operating loop instead of around it.

BuildWithHQ turns reusable platform primitives into products you can brand, extend, and scale.

buildwithhq.com / Build This