instinctools

How to agentize a business process

A 5-step playbook for shaping a first-pass agentic workflow, without spending weeks in discovery and thousands of dollars on consulting.

AI agents can summarize, classify, recommend, and act. But none of that guarantees that the business process around them will run any better.

74%

of leaders expect nearly half of their business processes to be **redesigned around AI agents within four years.**

Only 5%

of decision-makers say their business processes are **highly prepared for agentic AI today.**

1 playbook

to help you turn one business process into **a clear, actionable agentic workflow plan.**

Source: Deloitte, The path to agentic transformation, 2026

This playbook helps you do the first round of agentic process discovery yourself. Take one real workflow and work through five practical steps to determine:

- where the business value is
- what should be automated and what should not
- what AI needs to know and do
- where humans have to stay in control
- what it would take to run the process end to end

You do not need to select an LLM, map APIs, or design a production AI stack.

> Design the business process. The technology comes later.

By the end, you will have an Agentic Workflow Blueprint you can take into technical discovery.

STEP 1 OF 5

Pick a process worth fixing

Start with a business problem, not AI use

The wrong first question is:

> “ Where could we use an AI agent?

A better one is:

> “ Which process would benefit most from becoming faster, simpler, and more accurate?

Look for processes that happen often and which operational pain is both **visible and measurable** – start with **three to five** candidates. You will narrow them down to one process to take through the rest of the playbook.

How to prioritize your process candidates

Evaluate each process candidate/chosen process through three dimensions:

DIMENSION	WHAT TO CONSIDER	SCORE
01. Workload leverage	*How much work moves through this process?* Capture the actual volume: Nº ___ cases / requests / transactions / projects per time period	1 2 3 4 5 LOW HIGH **Low**: Relatively little workload or capacity is tied up in this process. **High**: Even a modest improvement could have a meaningful impact.
02. Operational friction	*How much unnecessary effort does the process create today?* Look for: manual work · delays · handoffs · chasing · rework · errors · duplicate entry · waiting for information	1 2 3 4 5 LOW HIGH **Low**: The process generally runs smoothly, with little unnecessary effort or delay. **High**: Frequent manual work, delays, rework, errors, or chasing create significant operational drag.
03. Business consequence	*What happens when the process moves slowly or goes wrong?* Look for: delayed revenue · customer experience · missed SLAs · operational bottlenecks · expensive rework · lost opportunities · compliance exposure	1 2 3 4 5 LOW HIGH **Low**: Problems have limited impact beyond local inconvenience or minor extra effort. **High**: Poor execution materially affects revenue, customers, SLAs, risk, or cost

The process with the **highest total score** is your strongest starting point: it combines meaningful workload leverage, operational friction, and business consequences when execution slows down or fails.

See what this looks like in practice

REAL-WORLD EXAMPLE

Ecommerce customer support

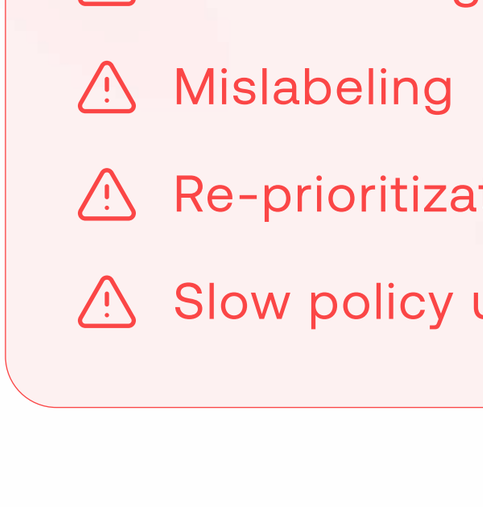
For one of our clients, a US ecommerce company, customer support was an obvious candidate: roughly 5,000 requests per day, with peaks up to 10,000, while manual processing took 6-12 minutes per ticket. The combination of volume, operational friction, and response-time pressure created a clear business case for redesigning the workflow rather than simply adding headcount.

CANDIDATE PROCESS	TYPICAL WORKLOAD	WORKLOAD LEVERAGE	OPERATIONAL FRICTION	BUSINESS CONSEQUENCE	TOTAL
Customer support request handling	~5,000 requests/day; up to 10,000 at peak	5*/5 - High leverage Thousands of requests move through the process every day.	5*/5 - High friction Manual ticket processing creates significant repetitive work at scale.	5*/5 - High impact Processing speed directly affects support capacity and customer response times.	15*/15

*Scores are illustrative. In your own assessment, use your company's actual operating priorities and thresholds.

- ⚠ Manual triage
- ⚠ Mislabeling
- ⚠ Re-prioritization
- ⚠ Slow policy updates





AI agents for customer support

See how Instinctools redesigned a support operation

Read case study →

YOUR TURN

Choose the process you will take forward

List 3-5 candidate processes. Capture the typical workload, then score each candidate from 1 to 5 using the same three dimensions.

CANDIDATE PROCESS	TYPICAL WORKLOAD	WORKLOAD LEVERAGE	OPERATIONAL FRICTION	BUSINESS CONSEQUENCE	TOTAL /15
1.		/5	/5	/5	/15
2.		/5	/5	/5	/15
3.		/5	/5	/5	/15
4.		/5	/5	/5	/15
5.		/5	/5	/5	/15

Not sure your shortlist is the right one?

If several process candidates look equally promising, or you want a second opinion before taking one forward, our **AI Adoption Workshop** can help you validate and prioritize them.

Together with Instinctools' experts, you will:

- ✓ compare opportunities against business value and feasibility;
- ✓ surface data, risk, and implementation constraints early;
- ✓ prioritize the strongest use cases and define clear next steps.

- Validate opportunities
- Assess risks & feasibility
- Get a clear roadmap

Explore the AI Adoption Workshop →

Before moving forward, check two things:

1. Can the outcome be measured?

☐ Yes ☐ No

You should be able to name at least one baseline KPI: cycle time, cost per case, throughput, error rate, manual touches, conversion, SLA attainment.

2. Can the risk be contained?

☐ Yes ☐ No

High-risk processes are not automatically off-limits. But consequential decisions should have clear boundaries, approval points, or human oversight.

If the highest-scoring process passes both checks, take it forward.
If it does not, move to the next-highest candidate.

STEP 1 OUTPUT – Process candidate brief

By the end of this step, you should be able to complete this sentence:

We are redesigning _____ process _____

_____

to improve _____ business outcome _____

_____ ,

currently measured by _____ baseline KPI _____

_____ .

NEXT →
Map the process and surface the friction

STEP 2 OF 5

Map the process and surface the friction

See how the work moves today, where people and systems get involved, and what keeps the process from running smoothly.

Before deciding where AI belongs, you need a clear view of the process **as it operates today**.

Reconstruct the current workflow, but don't try to document every click or edge case. Keep it to **5-10 major stages**, from the event that starts the process to the business outcome that completes it.

The goal is not to create perfect process documentation. It is to uncover where the workflow is losing **time, effort, or information**.

Start with:



For each major step, use the four questions below as friction lenses:

1. Who is involved?

Include anyone who performs work, provides input, reviews, or approves. Multiple roles or teams can indicate handoff and coordination friction.

2. What information and systems are needed?

List the information required to perform the step and where it comes from. Several sources, disconnected systems, missing data, or conflicting records can reveal information friction.

3. What is still done manually?

Describe what a person has to do by hand. Look for repetitive checking, copying, routing, data entry, follow-ups, or standard content preparation.

4. What typically slows, blocks, or sends the step backward?

Capture what goes wrong in practice. Think about missing inputs, approval queues, unclear ownership, failed checks, exceptions, and rework.

Then tag the friction you found:

- ☐ **Handoff**
  Work slows between people, teams, or approval points.
- ☐ **Information gap**
  Information is missing, fragmented, inconsistent, or difficult to access.
- ☐ **Manual work**
  Repetitive operational work still depends on a person.
- ☐ **Rework / exception**
  Work has to be repeated, corrected, or moved outside the standard path.

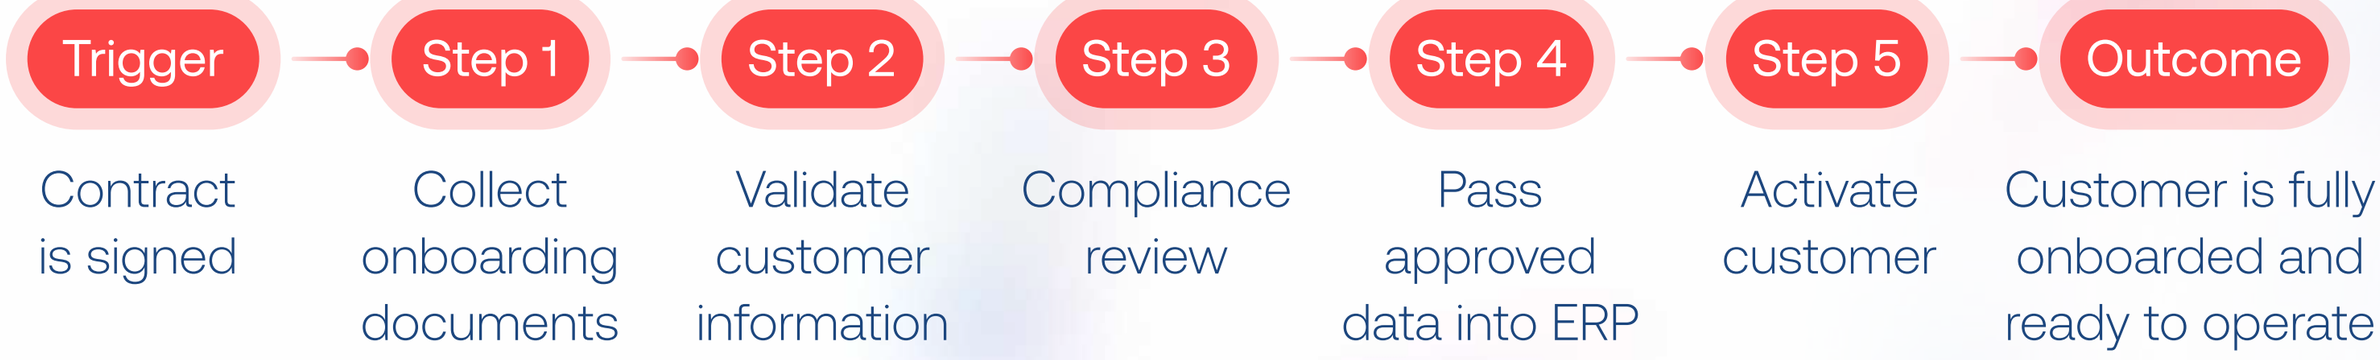
A stage can have more than one friction type. Do not look for AI opportunities yet. At this point, you are only diagnosing how the process performs today.

See what this looks like in practice

REAL-WORLD EXAMPLE

B2B customer onboarding

To see how the method works, consider a simplified B2B customer onboarding process:



PROCESS STAGE	WHO IS INVOLVED?	INFORMATION AND SYSTEMS NEEDED	WHAT IS DONE MANUALLY?	WHAT SLOWS, BLOCKS, OR SENDS THE STEP BACKWARD?	FRICTION
Collect onboarding documents	Sales Ops, customer	CRM, signed contract, documents received by email	Download attachments, check document list, follow up on missing items	Documents often arrive incomplete or through different email threads	☐ Handoff ☑ Information gap ☑ Manual work ☐ Rework / exception
Validate customer information	Sales Ops	CRM, registration documents, internal checklist	Compare company details across documents and CRM	Information does not always match across sources	☐ Handoff ☑ Information gap ☑ Manual work ☑ Rework / exception
Compliance review	Compliance, Sales Ops	Customer file, compliance policy, previous checks	Review case, request clarifications	Completed cases may wait in an approval queue	☑ Handoff ☐ Information gap ☐ Manual work ☐ Rework / exception
Pass approved data into ERP	Finance Ops	CRM, approved customer data, ERP	Re-enter customer information from CRM into ERP	Data-entry mistakes can send the case back for correction	☐ Handoff ☐ Information gap ☑ Manual work ☑ Rework / exception
Activate customer	Customer Ops, Finance, Sales	CRM, ERP, approval status	Check that previous steps are complete and notify stakeholders	Teams may not have one clear view of whether onboarding is complete	☑ Handoff ☑ Information gap ☐ Manual work ☐ Rework / exception

Without discussing AI yet, several problems are already visible:

- information is spread across multiple sources
- the same data is handled manually more than once
- approvals create handoffs and waiting
- mistakes push work backward
- no single system necessarily reflects the full process state

That is the level of insight you are looking for in your own workflow.

YOUR TURN

Map your process in one pass. Start with the first major stage after the trigger and continue until the business outcome is reached.

PROCESS STAGE	WHO IS INVOLVED?	INFORMATION AND SYSTEMS NEEDED	WHAT IS DONE MANUALLY?	WHAT SLOWS, BLOCKS, OR SENDS THE STEP BACKWARD?	FRICTION
1.					☐ Handoff ☐ Information gap ☐ Manual work ☐ Rework / exception
2.					☐ Handoff ☐ Information gap ☐ Manual work ☐ Rework / exception
3.					☐ Handoff ☐ Information gap ☐ Manual work ☐ Rework / exception
4.					☐ Handoff ☐ Information gap ☐ Manual work ☐ Rework / exception
5.					☐ Handoff ☐ Information gap ☐ Manual work ☐ Rework / exception

Now identify the friction that matters most

Your table may contain many inefficiencies. You do not need to solve all of them.

Choose the three friction points with the greatest impact on:

TIME · COST · CUSTOMER EXPERIENCE · QUALITY · EMPLOYEE EFFORT

№ 1

Friction: _____

_____

Why it matters: _____

_____

_____

_____

№ 2

Friction: _____

_____

Why it matters: _____

_____

_____

_____

№ 3

Friction: _____

_____

Why it matters: _____

_____

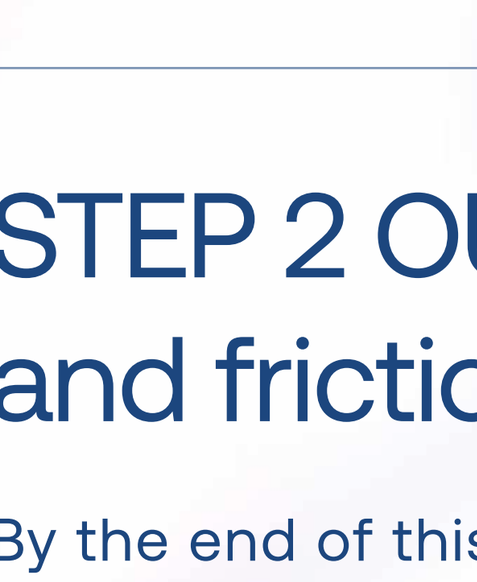
_____

_____

STEP 2 OUTPUT - Current-state process and friction map

By the end of this step, you should have:

- ✓ a clearly bounded process with 5-10 major stages
- ✓ visibility into the people, systems, and manual work involved
- ✓ three priority friction points worth addressing



AI sales representative

See how Instinctools connected fragmented sales tools into a living, self-updating context that an AI agent could use throughout the sales process.

Full case study →

NEXT →

Give every step to the right executor

STEP 3 OF 5

Give every step to the right executor

Not every friction point needs AI. And not every AI-powered workflow needs autonomy.

For some work, you will need conventional software. For other steps, AI can assist a person. Some activities may justify autonomous agent action. And some decisions should remain firmly human.

Execution models

RULE / SYSTEM

Deterministic logic executes the step based on predefined rules.

Choose this when the same inputs should reliably produce the same outcome.

AI ASSIST

AI helps a person perform the step, but does not complete the step itself.

Choose this when AI can interpret information, analyze, summarize, recommend, or prepare an output, while a person remains the **primary executor** of the step.

AI ACT

An AI agent itself completes the step or moves the workflow forward.

Choose this when AI needs to interpret changing context, make a bounded decision, use tools, create an output, update a system, or trigger the next stage of the process.

HUMAN

A person performs the core step.

Choose this when the core step depends on professional judgment, negotiation, relationship context, material consequences, or an exception that should not be delegated to AI.

Executor and oversight are two different decisions

First decide who performs the step: Rule/System, AI Assist, AI Act, or Human.
Then separately decide what human oversight is required:

NONE · REVIEW · APPROVAL · EXCEPTION ONLY

How to assign the right executor

Take each process stage from Step 2 and work through the questions below in order.

01 Can this step be completed reliably using fixed rules?

Examples: calculations, threshold checks, required-field validation, deterministic routing.

Yes → **RULE / SYSTEM**

Use deterministic automation when the logic is known.

No ↓

02 Does this step require interpreting, reasoning over, or making sense of variable information or context?

Examples: understanding an email, reviewing documents, comparing inconsistent information, case-specific circumstances.

No → **05 Is this step inherently human or non-digital?**

Examples: relationship-building, negotiation, physical inspection, leadership conversation, hands-on work.

No → **HUMAN**

Keep people where the work depends on human skills, expertise, or factors that should not be delegated to AI.

Yes → **AI ACT**

Let an agent execute when the work requires context and action — but can operate within clear boundaries.

Yes ↓

03 Can AI itself produce the required outcome of this step or move the workflow forward within defined boundaries?

Examples: create an output, classify and route a case, request missing information, create a plan, use tools, update a system, trigger the next stage.

No ↓

04 Can AI materially help a person perform this step?

Examples: analyze information, summarize documents, compare options, recommend, or draft a response.

Yes → **AI ASSIST**

Use AI to help a person perform the step — but it does not complete the step itself.

No → **CLARIFY / REDESIGN THE STEP**

AI has no clear role yet. Consider splitting the step, making the logic explicit, or standardizing it before introducing AI.

No → **HUMAN**

Keep people where the work depends on human skills, expertise, or factors that should not be delegated to AI.

Define human oversight (separately)

For any step where AI is involved, decide what level of human oversight is required.

None
No routine human intervention.

Review
A human checks the output before the process continues.

Approval
A human must approve the output before the process continues.

Exception only
A human gets involved only when defined conditions are met (e.g., low confidence, policy exception, risk threshold).

Keep in mind

- First decide who executes the step. Then decide the oversight level.
- Human review or approval does not automatically make a step "AI Assist."
- A step can combine an executor (e.g., AI Act) with human oversight (e.g., Review or Approval).

See what the execution model looks like in practice

REAL-WORLD EXAMPLE

Insurance partner onboarding

For one of our clients, a global insurance aggregator, onboarding a new carrier, MGA, or broker used to take 3-6 months. The process involved interpreting heterogeneous documentation, mapping API contracts, generating integration components, validating implementations, and navigating compliance requirements.

Instinctools redesigned the workflow with different stages being assigned to the executor best suited to the work: deterministic systems for hard checks, AI agents for context-heavy execution, and people for review, approval, and expert judgment.

The table below shows what that division of work looked like across the onboarding workflow.

PROCESS STAGE	PRIMARY EXECUTOR	HUMAN OVERSIGHT	WHY?
Documentation ingestion	AI ACT	Exception only / light review	Agent parses documents and extracts relevant material itself
Contract and reference analysis	AI ACT	Review	Agent builds reports and mines prior implementations; human performs a sanity check on large reports
Implementation plan	AI ACT	Approval / review	Agent scans the repository and produces the plan; developer approves or adjusts it for edge cases
Code generation	AI ACT	None routinely	Agent generates services, DTOs, mappings and tests
Build and self-fix	RULE / SYSTEM + AI ACT	None routinely	Deterministic checks identify failure; agent interprets errors and fixes the build
Smoke checks and PR creation	AI ACT	Approval	Agent runs checks and opens the PR; human reviews and merges
Compliance checks	AI ACT up to escalation	Human decision on exception	Agent detects/routes ambiguity; legal/SME owns the consequential judgment

YOUR TURN

Choose one primary executor for each process stage. If you select AI ACT, also check one human oversight option underneath it.

If a stage seems to need more than one executor, it may contain several different tasks. Split it into smaller steps and classify each separately.

PROCESS STAGE	RULE / SYSTEM	AI ASSIST	AI ACT	HUMAN
1.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐
2.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐
3.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐
4.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐
5.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐
6.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐
7.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐
8.	☐	☐	☐ AI Act **Human oversight:** ☐ None ☐ Review ☐ Approval ☐ Exception only	☐

STEP 3 OUTPUT - Execution model map

By the end of this step, you know:

- ✓ who should perform each major stage
- ✓ where AI should assist rather than act
- ✓ where AI can move the process forward itself
- ✓ and where human involvement remains necessary

NEXT →

Build the context AI needs

STEP 4 OF 5

Build the context AI needs

An agent can only act on what it can reliably understand about the business, the case, and the current state of work.

In Step 3, you identified where AI could assist people and where an agent could execute work itself.

Now focus only on those **AI Assist** and **AI Act** stages.

For each one, ask:

> “ What would AI need to know to perform this step correctly – the way a capable employee would?

Do not think about APIs, RAG, vector databases, or memory architecture yet. At this stage, your job is to identify the business context the workflow requires and where that context lives today.

Context can come from three sources

Business knowledge – how the company works.

Examples:

* policies and SOPs;
* service catalogs;
* product rules;
* pricing or eligibility logic;
* compliance requirements;
* internal terminology;
* approved templates and guidelines.

Case context – what is true about this specific situation.

Examples:

* customer profile and history;
* emails and conversations;
* contracts and submitted documents;
* transactions;
* previous interactions;
* existing project or matter records;
* prior decisions.

Process state – what has already happened and what happens next.

Examples:

* which stages are already complete;
* what is still missing;
* who has approved what;
* which task is open;
* what exception has been raised;
* what the next expected action is;
* whether the workflow is waiting for a person, system, or external party.

Knowing the data is not the same as knowing the process state

An agent may know who the customer is and still have no idea what should happen next. That is why complex agentic workflows need a shared operating layer that keeps business context, case context, and process state connected as work moves across agents, systems, and people.

How to map the context

Take each AI Assist or AI Act stage from Step 3 and answer four questions:

1. What does AI need to know?

Be specific. Instead of “customer data”, write: “Current opportunity stage, previous conversations, decision-makers, and open follow-up tasks.”

2. Where does that information live today?

Think in business terms:

CRM · ERP · EMAIL · MESSENGERS · DMS · SHAREPOINT · JIRA · SPREADSHEETS · INTERNAL APPLICATIONS · EMPLOYEE KNOWLEDGE

You do not need to know how those sources will technically connect.

3. Can you rely on it?

Ask whether the information is:

- [ ] Current and accessible
- [ ] Fragmented or inconsistent
- [ ] Incomplete / missing

This is important because giving an agent access to five systems does not automatically give it reliable context.

See what this looks like in practice

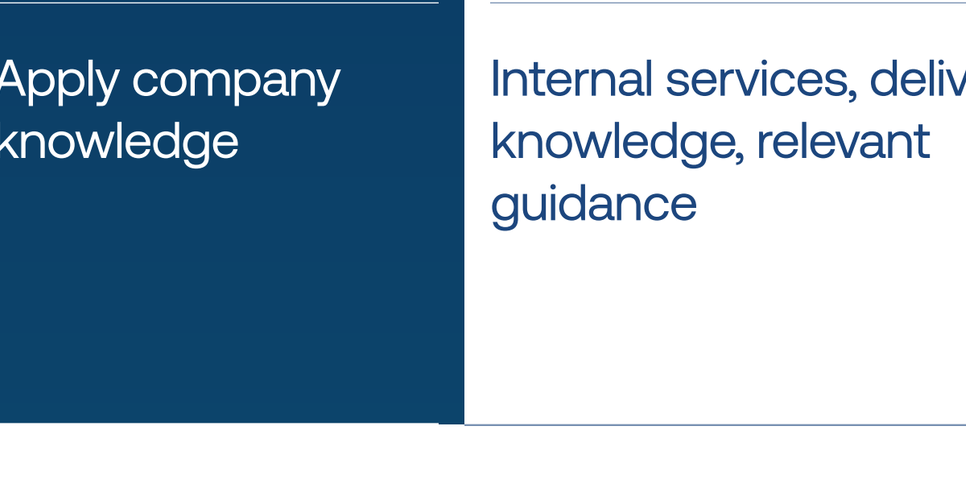
Choose one primary executor for each process stage. If you select AI ACT, also check one human oversight option underneath it.

For one of our clients, a consulting company, sales representatives worked across a highly fragmented toolset: HubSpot for deal records, Teams for conversations, OneDrive for documents, Jira for project tracking, Power BI for reporting, Miro for collaboration, and LinkedIn for prospecting.

The company’s first agents struggled because their context was **static and manually maintained**. Employees still had to collect updates from multiple applications and feed them back into the AI.

Instinctools connected the agent across the company’s working environment so it could maintain a **living, self-updating project context and use that context** to suggest or take next steps throughout the sales pipeline.

AI-ENABLED STAGE	WHAT AI NEEDS TO KNOW	WHERE THE CONTEXT LIVES TODAY	CONTEXT STATUS
Assess an active opportunity	Customer, contacts, opportunity stage, previous deal activity	HubSpot	☑ Current and accessible
Understand the latest conversation	Recent questions, commitments, objections, decisions	Microsoft Teams	☑ Fragmented or inconsistent
Determine what is still outstanding	Open tasks, dependencies, unresolved actions	Microsoft Teams + Jira	☑ Fragmented or inconsistent
Use relevant project materials	Proposals, requirements, shared documents	OneDrive	☑ Current and accessible
Recommend the next step	Current pipeline state + previous actions + outstanding work	Multiple systems	☑ Fragmented or inconsistent
Apply company knowledge	Internal services, delivery knowledge, relevant guidance	Confluence	☑ Current and accessible



AI agents for partner onboarding

See how Instinctools helped a global insurance aggregator redesign partner onboarding from a months-long process into a governed multi-agent workflow.

Full case study →

YOUR TURN

Go back to Step 3 and copy only the process stages you classified as **AI Assist** or **AI Act**.

For each stage, identify what AI would need to know to perform the work correctly.

PROCESS STAGE	WHAT DOES AI NEED TO KNOW?	WHERE DOES THE CONTEXT LIVE TODAY?	CONTEXT STATUS
1.			☐ Current and accessible ☐ Fragmented or inconsistent ☐ Incomplete / missing
2.			☐ Current and accessible ☐ Fragmented or inconsistent ☐ Incomplete / missing
3.			☐ Current and accessible ☐ Fragmented or inconsistent ☐ Incomplete / missing
4.			☐ Current and accessible ☐ Fragmented or inconsistent ☐ Incomplete / missing
5.			☐ Current and accessible ☐ Fragmented or inconsistent ☐ Incomplete / missing
6.			☐ Current and accessible ☐ Fragmented or inconsistent ☐ Incomplete / missing

What does your context map tell you?

Now scan the table for patterns. For every question you answer “yes” to, the note underneath shows what that means for your workflow.

- **Is the same context needed across several stages?**
  - → **Shared context may be required**. The workflow may need a common context layer rather than each agent retrieving and maintaining information independently.
- **Does one stage depend on several systems?**
  - → **Integration complexity increases**. The more sources a stage depends on, the greater the risk of fragmented, inconsistent, or outdated context.
- **Is critical information stored in conversations, documents, or people’s heads?**
  - → **Context preparation may be needed first**. Identify what knowledge must be captured, which source should be authoritative, and how AI will reliably access the latest version.
- **Does AI need to know what happened earlier in the workflow?**
  - → **Process continuity matters**. AI needs access not only to information, but also to the decisions, actions, and state carried forward from previous stages.
- **Do different sources disagree?**
  - → **A source-of-truth policy is needed**. Define which source takes precedence and how conflicting information should be resolved.

Watch for these to see where context gets complicated

- [ ] AI would need to search across several systems for one decision
- [ ] The same information exists in several places
- [ ] Nobody is sure which source is authoritative
- [ ] Context changes while the process is running
- [ ] Later stages depend on decisions made earlier
- [ ] Different agents would need access to the same information

The more boxes you check, the more likely you are to need a shared **operating layer** where AI agents and systems can work with the same context, carry state forward, and operate safely across the workflow.

STEP 4 OUTPUT - Context map

By the end of this step, you should know:

- ✓ what each AI-enabled stage needs to understand
- ✓ which information is stable and which changes case by case
- ✓ where that context lives today
- ✓ whether the process needs continuity from one stage to the next
- ✓ where missing or fragmented context could prevent reliable execution

NEXT →

Decide how the workflow should be orchestrated

STEP 5 of 5

Decide how the workflow should be orchestrated

You now know how the process works, where AI belongs, where people stay in control, and what context AI needs to perform reliably.

Now step back and look at the workflow as a whole.

What needs to keep all of these moving parts coordinated from one stage to the next?

As work moves through the process, something needs to keep the **current process state, business context, systems, AI capabilities, and human controls aligned**.

This is the role of an **agentic operating system**, which gives your agents, applications, and people a shared environment in which the workflow can run reliably from beginning to end.



See what this looks like in practice

REAL-WORLD EXAMPLE

Solving an operating-layer problem across individual AI solutions

A German professional services group wanted to introduce AI across legal, tax, assurance, and consulting operations. There was no single “agent problem.” Different parts of the business needed different capabilities:

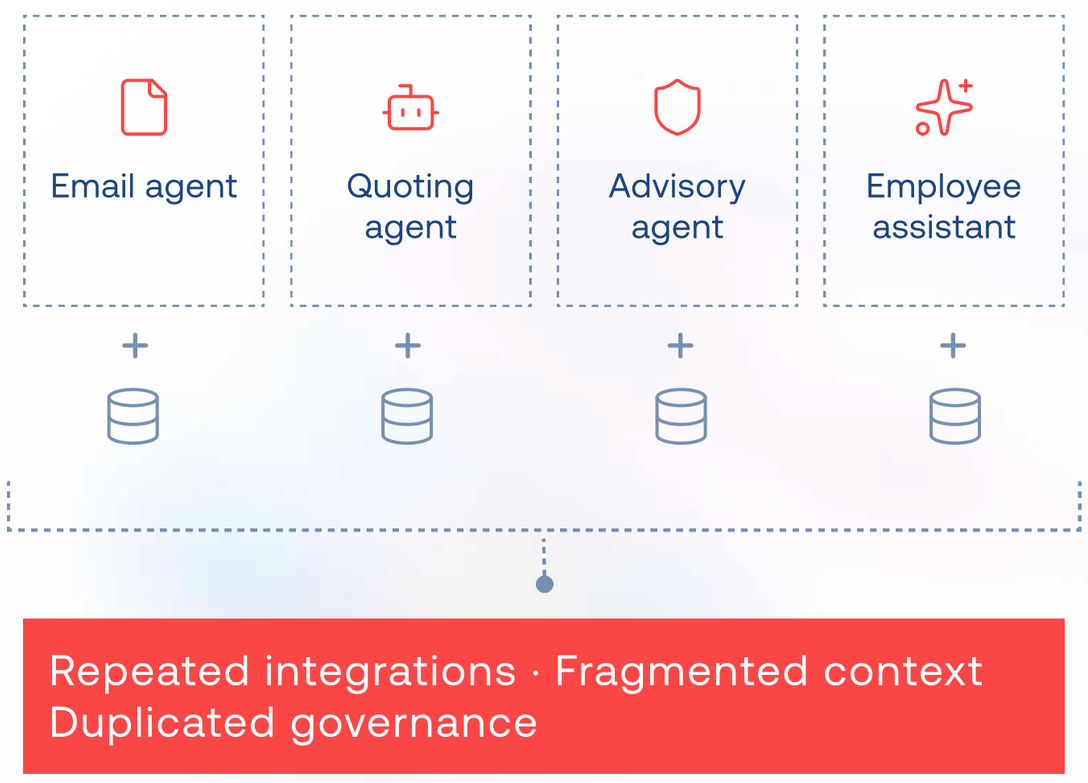
* an email intake agent to classify and route requests;
* a quoting agent to prepare fee proposals;
* an advisory agent to analyze complex documents;
* an employee assistant for routine internal questions.

Building these as unrelated AI tools would have solved individual tasks, but it would also have recreated the fragmentation the company was trying to remove.

So Instinctools built the **operational foundation**: shared integrations, customer and matter context, orchestration, data access, governance, and auditability. Individual AI capabilities could then operate on top of that common layer and reuse it across business units.

Without a shared operating layer

Individual solutions lead to duplicated effort, fragmented context, and limited visibility.



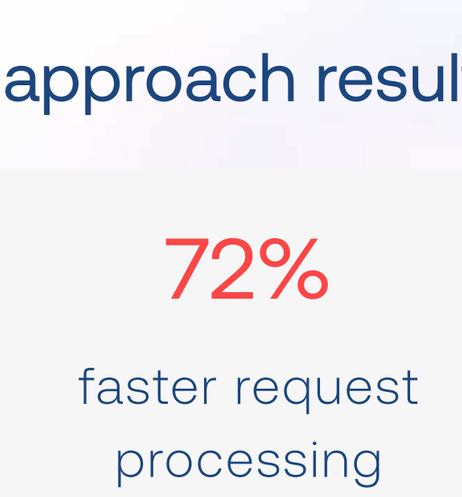
With a shared operating layer

Different capabilities work together on a common foundation.



This approach resulted in:

72%	65%	~6 hours
faster request processing	less fee proposal preparation time	of repetitive admin work saved per week



AI operational platform for a professional services firm

See how Instinctools replaced a collection of isolated AI initiatives with a shared operational foundation for context, integrations, orchestration, governance, and reuse across business units.

Full case study →

YOUR TURN

Assess your orchestration needs

Look back at the workflow you designed in Steps 1-4. Check every statement that applies.

Across the process

- [ ] The workflow has several stages that depend on one another
- [ ] Work moves between AI, deterministic systems, and people
- [ ] The process may continue over hours, days, or longer
- [ ] What happens next depends on the outcome of an earlier stage

Across context

- [ ] Context created or updated in one stage is needed later
- [ ] Several AI-enabled stages need access to the same business context
- [ ] Context changes while the workflow is running
- [ ] Information comes from multiple sources that need to remain consistent

Across systems

- [ ] The process crosses multiple enterprise systems
- [ ] Different stages happen in different tools or applications
- [ ] Agents need to work with existing CRM, ERP, collaboration, document, or legacy systems
- [ ] A change in one system needs to be reflected elsewhere in the process

Across people and control

- [ ] Human reviews or approvals occur at specific stages
- [ ] People step in when predefined exceptions occur
- [ ] Different roles have different authority or access
- [ ] AI actions and human decisions need to be traceable
- [ ] Managers need visibility into where the process stands, what is blocked, and who or what needs to act next

The more boxes you check, the more orchestration the workflow requires.

Give your AI agents the structure they need to work safely inside production

Once context, systems, AI actions, and human decisions become interdependent, the real challenge is no longer building an agent, but keeping the entire process coordinated from one stage to the next.

GENiE is Instinctools' business agentic operating system that brings AI agents, enterprise systems, business context, rules, and human control together into one coordinated workflow.



Connects context across stages

So agents and people work from the same current, governed view of the business.

Orchestrates AI, systems, and people

So work moves end to end without breaking at handoffs.

Keeps execution visible and under control

So teams can spot blockers, deviations, and risks before they derail the process.

SCALE AGENTIC AI WITHOUT LOSING CONTROL

Overcome fragmented context, costly handoffs, and reactive execution to turn isolated AI wins into enterprise-wide value.

Explore GENiE →

STEP 5 OUTPUT - Your agentic workflow blueprint 🎉

You now have enough information to describe your workflow as a system.

You have defined:

01 **The opportunity**
What business case you want to improve.

02 **The workflow**
How work moves today and where the friction sits.

03 **The execution model**
Which stages belong to rules, AI assistance, AI action, or people and where human oversight remains.

04 **The context**
What AI needs to understand, where that information lives, and where continuity or fragmentation becomes a challenge.

05 **The orchestration needs**
How to keep all moving parts connected and coordinated as work progresses end to end.

You now have a clear starting point

It is not an implementation specification yet. But it is enough to move from broad AI exploration to a concrete view of the process, where AI belongs, what it needs to work effectively, and how to make the workflow run reliably end to end.

Ready to turn your blueprint into an implementation plan?

Pressure-test your assumptions for technical feasibility and operational safety before committing to a build.

Validate your workflow with Instinctools' AI architect →

Practical AI. Real business impact

instinctools

We advise, engineer, and automate with innovation in mind



“Great ideas don’t grow alone — teamwork transforms challenges into breakthroughs, especially in the age of AI, where possibilities redefine innovation.”

Alexey Spas
CEO and Co-founder

- Team augmentation and scaling
- Digital product engineering and AI enablement
- Business process automation and AI transformation
- Data infrastructure and analytics
- Cloud infrastructure and DevOps

25+
Years in software engineering

400+
Professionals worldwide

650+
Completed projects

CRÉDIT D'IMPÔT RECHERCHE AGRÉMENT CIR

CERTIFIED ISO 27001 2022

CERTIFIED ISO 9001 2015

CERTIFIED ISO 37001 2016

CERTIFIED ISO 14001 2015

CERTIFIED ISO 45001 2018

USA

12430 Park Potomac Ave, Unit 122
Potomac MD 20854, USA (GMT–5)

Germany

Hauptstaetter Str. 89,
70178 Stuttgart, Germany (GMT +2)

Development centers

Poland (Warsaw)
India (New Delhi)
Kazakhstan (Astana)

instinctools.com

contact@instinctools.com

Copyright 2000–2026