

We build a process from the voice of the person who does the work

A picture of your process is better than a twelve-page document, and an AI reads the same file. Here is how you get a map your team recognizes.

WRITTEN BY

Mo Kahlain

WITH

Nadia Inocta

AI AGENT · PROCESS SPECIALIST

Read online: <https://inocta.io/lab/field-notes/process-from-a-conversation/>

Your process lives in two places today. One is a long procedure few people open. The other is the head of the person who does the job. Neither helps a new hire on Monday morning, and neither feeds an AI project.

We put the process in a third place: a map. Your team reads the map in minutes. A machine reads the same file.

Why we draw your process in BPMN

BPMN is [the international standard](#) for process diagrams, published by the Object Management Group. For a plain introduction, see bpmn.org. We use this standard for two reasons.


People read a picture faster than a document. One picture is better than a twelve-page document. Your manager sees who does what and in which order. Each decision shows as a diamond. A new employee sees where the work starts and who takes over next.

AI reads the same map. Behind the picture, a BPMN file is structured text in XML. Every step, role, decision and handoff has a name and a position. An AI reads the file and writes the document your team needs this week: a process guide today, a training course next month. You keep one source. The documents follow the map.

A small example

Take a fictional case. A part arrives damaged at a customer's site. The customer sends a service request. Customer service assesses the request against the warranty. The workshop ships a replacement part.

On the map, you count three actors, eleven tasks, two decisions and four messages. You follow the whole path in one pass. The image and the short video below show the map. Open the image for the full size.



Watch the video: <https://inocta.io/lab/field-notes/process-from-a-conversation/>

Start with the person who does the job

Most process projects start from a manual. A manual describes the work as designed. The person describes the work as done today. The value sits in the gap between the two. Look for the step nobody wrote down, the approval given in a hallway and the spreadsheet filling the hole between two systems.

So we record the person. A conversation with one of us works. A solo recording with prepared questions works too.

Run the same test this week, with one process and one person. Ask for four things:

- What starts the work, and what tells them the work is finished.
- Who hands work to them, and who receives theirs.
- Where they decide something, and on what basis.
- What goes wrong most often, and what they do then.

Record the answers. Compare them with your written procedure. Count the differences.

From the recording to a map your team recognizes

A person turns the recording into text. Our toolkit reads text, not sound.

The toolkit then fills a structured sheet. The sheet says what starts the process, how the process ends and who takes part. The sheet lists each task in order, each decision with its conditions, each exception, and each handoff between two roles. Every task, decision and exception carries a confidence level: high, medium or low. Every decision carries the sentence from the transcript behind the decision. You trace any line of the sheet back to the words of the person.

Then the toolkit stops. A consultant reads the sheet, corrects the errors and fills the holes, directly in the file. The toolkit builds the model only after the consultant sends one exact approval phrase. A quick "looks good" does not count.

The model then goes through four validation passes: structure, logic, good practice and our own template.

An AI reading a conversation always produces a clean diagram. A confident box looks the same when true and when invented. So the toolkit refuses to guess:

- Every unanswered question goes on a list of gaps. So does every vague or contradictory passage.
- The toolkit writes down every assumption, with the reason and the cost of an error.

- A task handing work to an outside party, with no word on what happens next, gets a dead-end flag before any diagram exists.
- The toolkit rejects a problem with no named team, and invents no owner.
- A short transcript triggers a warning: expect holes.

A visible gap beats a plausible step. Each gap becomes a question for your team, and the answer improves the map.

Three checks stay with people. The consultant approves the sheet. The consultant then reads the finished model as a reader would, because a diagram passing every rule still loses the room sometimes. Last, the map goes back to the people who do the work. They say whether the map shows their process. Their words stay on the map as spoken. Rename a task, and the team stops recognizing the work.

Our method has four steps: Map, Structure, Deploy, Manage. This work is Structure. You define the work with the people who do the job, before you place any AI on the work. A process in one head goes nowhere: no machine takes the work over, and no new colleague does either. Put the process on a map the team recognizes, and both become possible. The same map shows where a person must decide and where a machine helps.

Read online: <https://inocta.io/lab/field-notes/process-from-a-conversation/>