

FREE GUIDE

The One-Hour Scenario Assistant

AI writes the scenario. You make every design call.

Writing one good practice scenario used to take me days. Now it takes me 30 to 60 minutes. Not because AI does the thinking, but because it does the typing. The One-Hour Scenario Assistant is a simplified version of the process I use.

You bring the course content: a policy, a manual, a procedure, a stack of documents. The assistant takes you in four steps to one practice scenario with three decisions, realistic consequences and feedback that shows instead of lectures. At four moments it stops and asks you to decide. Those decisions are what make a scenario good, and they're the part AI can't make for you.

Step	The assistant drafts	You decide	Time
1. Find the behaviour	Actions hidden in the content	The one action this scenario trains	10 min
2. Find the mistakes	Likely mistakes and their causes	Which mistakes are real	10 min
3. Build the skeleton	Character, challenge, ideal path, choices	Whether the wrong choices are tempting enough	15 min
4. Write the scenes	Situations, choices, feedback	What sounds real, and what the expert must check	20 min

What's in the kit: the assistant (a text file and a skill version), this guide with a start prompt, a real worked example, a clickable version of that scenario, practice content, and a final checklist.

THE NEXT STEP

This kit builds one scenario. In my online programme, the **Learning Lab**, you build your own AI design process for every project. The waitlist is free and gets first pick of places and the lowest founding prices. **More on the last page.**

Set it up (5 min)

The assistant is a set of instructions that your AI tool follows through the whole process. It works in ChatGPT, Claude, Gemini and Microsoft Copilot. There are two ways to use it.

Quick start, in any tool. Open a new chat. Paste the full contents of *one-hour-scenario-assistant.txt* as your first message, with this line above it: *Follow the instructions below for this whole conversation.* Then send your start prompt (next section).

To use it again and again, save it once:

Tool	Save it for repeated use
ChatGPT	Add it as a skill using the <i>SKILL.md</i> file in the kit, or paste the text file into a project's instructions.
Claude	Add it as a skill using <i>SKILL.md</i> , or paste the text file into a project's instructions.
Gemini	Add it as a skill using <i>SKILL.md</i> . Skills replace Gems and may need a paid plan; the quick start always works.
Microsoft Copilot	Create an agent and paste the text file into the instructions field.

AI tools rename their menus often. If a button has moved, look for "skills" or "instructions" in the settings. Not allowed to put client material into AI? Use the practice content at the back of this guide, or something already public, such as a published policy.

Your first message: the start prompt

The assistant works with whatever you give it, but the more it knows about your learners, the less it has to guess. Copy this into your first message and fill in what you know. Anything you leave open, the assistant asks you in step 1.

I want to build a practice scenario.

Learners: [their role, how experienced they are, where they work]
Language of the scenario: [for example English or Dutch]

What I know so far:

- The business problem: [...]
 - What people do now that isn't working: [...]
 - What success would look like: [...]
- (Don't know yet? Leave it open and ask me.)

Here's the content:
[paste your content]

Start with step 1.

Three tips for better input. Describe the learners as concretely as you would to a colleague:

"customer service agents, most in their first year, working from home" gives a far better scenario than "employees". Paste text, not a link: most AI tools can't open files behind a login. And if you only know

the problem in your own words ("people keep promising refunds they can't give"), write it down just like that. That's often the best starting point.

No content at hand, or not allowed to use client material? Use the practice content at the back of this guide. You can try the whole process in under an hour, without anything from a client.

Step 1: Find the behaviour (10 min)

Content tells you what people should *know*. A scenario needs what people should *do*. So the assistant starts by finding the one action on the job that this e-learning has to change.

What the assistant does. It asks three questions: what business problem the e-learning should solve, what people do now that isn't working, and what success would look like. It helps you turn the answers into a goal: *[metric] will [increase/decrease] by [amount] by [date] as people [do the action]*. If your goal is vague, it asks what you'd see if you filmed someone doing it well. Then it lists the actions in your content that pass that test, and the content that supports no action at all.

Your call. Pick *one* action for this scenario. Choose the one where mistakes are costly and common, not the one with the most content behind it. Then apply the photo test: if you filmed someone doing this well, what would you see?

Why you and not AI? Because the choice depends on what isn't in the content: what actually goes wrong at work, what the requester worries about, what learners already do well.

Step 2: Find the mistakes (10 min)

The wrong choices in your scenario are only as good as the mistakes behind them. Real mistakes make a scenario feel true. Invented ones turn it into a quiz with an obvious answer.

What the assistant does. For the action you chose, it lists likely mistakes: what people do, why it feels reasonable in the moment, and what happens next. It adds two almost-right actions, and marks what comes from your content and what it assumes. Because the real mistakes are rarely written down anywhere, it also gives you four or five questions to take to an expert or an experienced employee: where do people hesitate, which two options both sound defensible, what do beginners do that experienced people don't.

Your call. Show the list to one person who does this work, or to the expert. Ask one question: *"Which of these have you actually seen?"* Keep two real mistakes and one almost-right action. They become your bad and OK choices in step 3.

Why you and not AI? AI produces mistakes that sound plausible. Only practice tells you which ones people make. A mistake nobody makes becomes a choice nobody picks.

Step 3: Build the skeleton (15 min)

Start from the best path and build backwards. Once you know what doing it well looks like, the wrong turns are easy to place. Three decisions is enough for a first scenario.

What the assistant does. It proposes a character in your learners' own role with a natural goal for the day, the person they deal with, a setting they'll recognise, the moment that starts the scenario, and what's at stake. Then three decisions that escalate, each with a good, an OK and a bad choice built on the mistakes you kept.

Your call. Cover the labels and read the choices as a learner would. Would a tired colleague on a busy Friday pick the bad one? If the good choice gives itself away (the longest, the most careful-sounding, always first), rewrite it until all three feel reasonable.

Why you and not AI? AI tends to make the right answer sound right. Making the wrong answer tempting takes someone who knows why people choose it.

Step 4: Write the scenes (20 min)

Now the skeleton becomes something a learner steps into. The difference between a scene and a description of a scene is one rule: show, don't tell.

What the assistant does. It builds each situation as context, trigger, observable behaviour, tension and choice. It shows what the character sees and hears instead of naming conclusions, cuts interpretive words like "busy" or "stressed", and never lets the narrator point out a mistake. Every choice is a complete sentence about what the learner does or says, similar in length to the others, and the order is mixed, so the right one doesn't give itself away. Feedback names the effect instead of a verdict and explains it through the other character's eyes. You get a complete scenario: an intro that puts the learner in the role, the three decisions, and an outro with three endings and a reflection question. Plus a list of facts your expert should check.

Your call. Read the scenario out loud. Anything that sounds like a policy document instead of a person gets rewritten. Watch for words like *busy*, *angry* or *risky*: they draw the conclusion for the learner. For every sentence, ask what you would see or hear if you were standing next to this person, and write that instead. Then mark every factual claim the expert must check before it goes live, and run the final checklist at the end of this guide.

WORKED EXAMPLE

From delivery procedure to customer call

A real run of the assistant in ChatGPT. The shop, the names and the numbers are made up; the AI output and the corrections are real.

The input. A one-page procedure for handling complaints about late deliveries at HomeNest, an online furniture shop, plus the filled-in start prompt. Both are at the back of this guide, so you can run the same content yourself and compare.

Step 1. Goal: *negative reviews about late deliveries drop from 31% to 20% by June 2027, as agents only give confirmed delivery dates.* The action we chose: *give the customer a new delivery date only when it's confirmed in the system, and never guess a date.*

Step 2. The assistant listed seven likely mistakes and two almost-right actions. A team lead recognised two: the hedged date ("It should probably arrive Friday, but I can't promise") and passing on the carrier's "expected tomorrow" as if it were confirmed. The assistant also found a gap: the procedure never says what an agent should do when there is no confirmed date. That became the good choice: promise a concrete callback.

Step 3. You're Jamie, in your first year on the customer service team. Eva has taken the day off to wait for her dining table. It didn't come.

Decision	The moment	Best choice
1	Eva: "Is it coming tomorrow or not?" The carrier note says "expected tomorrow"	Say the date isn't confirmed, so she shouldn't plan around it
2	Eva: "Just give me your best guess." The call timer passes 5:20	Promise a callback from planning before 4 p.m.
3	3:52 p.m. A planner writes: "Probably Friday, I'll enter it tomorrow"	Call Eva before 4, say Friday is expected but not confirmed, agree a time tomorrow

Step 4. The full scenario, after the corrections on the next page. Want to click through it? Open [homenest-scenario-demo.html](#) from the kit in your browser.

Intro

It is 2:35 p.m. at the HomeNest customer service desk. You're Jamie, in your first year on the HomeNest customer service team. You have Eva's order open in the CRM while the call timer runs in the corner of your screen. The queue counter shows 7 callers waiting. This morning, your team lead Priya said the team needs to bring average call time down to 6 minutes.

Decision 1

Eva says, "I've been home all day. Is it coming tomorrow or not?"

You look at the CRM. The carrier note says "expected tomorrow," but there is no confirmed delivery date in the order.

Choice	Feedback label	Why
"The carrier expects it tomorrow, so that's the date I'd plan for. I'll let you know if anything changes."	Tomorrow goes in the diary	Eva says, "Fine, I'll tell work I need tomorrow off too." You hear keyboard clicks on the other end of the line.
"The carrier expects it tomorrow, but that isn't a confirmed date yet, so I don't want you to plan around it."	The diary stays open	Eva pauses. "Okay. I won't book tomorrow off yet. What happens next?"
"I don't have a confirmed date yet. I'll send you a message if the CRM updates later today."	Waiting without a time	Eva says, "So I just keep checking my phone?" A few seconds later she adds, "I need to tell my manager before I leave today."

Decision 2

Eva says, "Just give me your best guess. I need to know whether to take another day off."

You refresh the order. The date field is still blank. The call timer passes 5:20. Above it, the queue counter still shows 7.

Choice	Feedback label	Why
"Friday looks most likely from what I can see, but I can't promise it. That's my best estimate."	Friday becomes the plan	Eva says, "Right, I'll ask for Friday off then." She repeats the date once more before ending the call.
"I can't confirm a date yet. I'd wait until tomorrow morning before deciding whether to take the day off."	Decision pushed forward	Eva says, "My manager wants an answer today." She asks whether there is anyone else she can call herself.
"I can't give you a date yet. I'll check with planning and call you back before 4 p.m."	A time to expect	Eva says, "Before four? Okay, I can wait for that." She gives you the best number to call.

Decision 3

It is 3:52 p.m. You have Eva's order open on one screen and Teams on the other.

A planner has just replied: "Probably Friday, I'll enter it in the system tomorrow morning."

The CRM still shows no confirmed date. You promised Eva a call before 4 p.m.

Choice	Feedback label	Why
I'll wait until planning enters Friday in the CRM tomorrow, then call Eva with the confirmed date, even though it's after 4.	Four o'clock passes	At 4:06, Eva looks at her phone, then at the clock. She opens the HomeNest chat and types, "Jamie said someone would call me before four."

Choice	Feedback label	Why
I'll call Eva before 4 and tell her planning says Friday. I'll explain that the CRM will be updated tomorrow morning.	Friday sounds settled	Eva says, "Great, then I'll take Friday off." Before you finish explaining the CRM update, she opens her work chat and types, "I need Friday off for the delivery."
I'll call Eva now, say Friday is expected but not confirmed, and agree a time tomorrow morning when I'll call with the confirmed date.	Friday stays pencilled in	Eva says, "Okay, I won't ask for Friday off yet." She opens her calendar and says, "Call me by 10 tomorrow morning. I'll wait until then."

Outro

The next morning, you open Eva's order before the agreed callback time.

Strong: Eva says, "Thanks. I waited until you had the confirmed date before I changed anything at work."

Mixed: Eva says, "I kept checking my phone because I wasn't sure when I'd know for certain."

Poor: Eva says, "I rearranged work because I thought the date you mentioned was the one I should plan around."

You refresh the order. The delivery date field now reads Confirmed: Friday, 18 June 2027.

Reflection: when a customer pushes you for a date that is not confirmed, what could you say that gives them something concrete without turning an estimate into a promise?

What the designer changed

The first draft was a good start, not a finished scenario. These are the corrections made at the checkpoints, in this real run:

AI draft	What we changed	Why
The good choice was the longest and most careful-sounding in all three decisions	Made all three choices similar in length and tone	Learners pick the longest answer when unsure
The OK choice was "contact us again" three times	A different almost-right action in every decision	A repeated pattern teaches the pattern, not the skill
Decisions 2 and 3 were the same dilemma	A new decision 3: the planner's "probably Friday" in Teams	Each decision should test a different mistake
The story talked about "Jamie", the choices said "I"	The learner is "you" throughout	The learner should be in the role, not watching it
No pressure anywhere	A queue counter, a call timer and a 6-minute target	Pressure is what makes the wrong choice tempting
Closing lines summed up the learner's performance and assumed one route	Eva's own words, true for every route, and a concrete date	Show the result; don't grade it

None of these changes took long. All of them needed someone who knows how learners choose. The assistant now has each of these corrections built in as a rule, but the checkpoints are still where you catch what it misses.

FINAL CHECKLIST

Before your scenario goes live

Cover the feedback labels and read the scenario as a learner would. Then check:

- ☐ Every choice is something the learner says or does, not something they know.

- ☐ Each wrong choice is a mistake people actually make, confirmed by someone who does the work.

- ☐ The good choice isn't the longest, the most careful-sounding, or always in the same place.

- ☐ Each situation starts in the middle of the action and shows only what the character sees and hears.

- ☐ Each situation follows from the one before.

- ☐ Feedback names an effect and shows what happened, instead of giving a verdict or a lecture.

- ☐ No interpretive words like busy, angry or risky: every conclusion can be drawn from what is shown.

- ☐ The names, setting and language match your learners.

- ☐ The outro has three closing lines (strong, mixed, poor) and one reflection question.

- ☐ The expert has checked every factual claim you marked.

Stuck on one of these? Go back to the step it belongs to and ask the assistant for a new version. You don't have to start over.

Try it yourself, with the same content

This is the fictional procedure used in the worked example. Paste it into the assistant with the start prompt below, make your own calls at every checkpoint, and then compare your scenario with ours. You won't get the same result: AI never writes the same thing twice, and your decisions will differ from ours.

Handling complaints about late deliveries
Procedure for customer service agents, HomeNest online furniture

About HomeNest

HomeNest was founded in 2011 and now delivers furniture to 40,000 customers a year across the Netherlands and Belgium. Our values are care, clarity and reliability.

When a delivery is late

Most complaints about late deliveries come in by phone or chat on the day the delivery was planned. Customers have often taken a day off work to wait for it.

1. Check the order in the CRM before you promise anything. Look at the delivery status and the notes from the carrier.
2. Acknowledge the problem first. Do not blame the carrier or the warehouse, even if they caused the delay.
3. Give the customer a new delivery date only when it is confirmed in the system. Never guess a date.
4. You may offer compensation up to EUR 25 (a voucher or a refund of the delivery costs) without approval. For anything above EUR 25, ask your team lead first.
5. Record the complaint in the CRM with the category "Late delivery" and a short summary of what you agreed.
6. If the delivery is late for the second time, escalate to the planning team the same day. Do not just give a new date.

Tone of voice

Use the customer's name. Keep sentences short. Avoid internal terms such as "slot", "route" or "hub".

Facts and figures

In 2025, 6% of all deliveries were late. Late deliveries cause 31% of all negative reviews. A customer who receives a confirmed new date within the first call is twice as likely to order again.

The start prompt, filled in for this content

I want to build a practice scenario.

Learners: customer service agents at an online furniture shop, most in their first year, working by phone and chat

Language of the scenario: English

What I know so far:

- The business problem: late deliveries cause a third of our negative reviews
- What people do now that isn't working: agents guess a new delivery date to calm the customer down, and the date is often wrong again
- What success would look like: fewer negative reviews about late deliveries

Here's the content:

[paste the HomeNest procedure]

Start with step 1.

What to look out for. Does the good choice give itself away by being the longest? Is the learner "you" throughout? Can you see the pressure? Those were the points we had to correct in our run.

WHERE TO GO FROM HERE

One scenario is a start. Your own process is the goal.

The real gain comes when this becomes *your* way of working, and when it gets a little better with every project: you notice which instructions you keep correcting, which checks you skip, which calls you want to keep making yourself.

That's what the **Learning Lab** is for: my online programme for experienced instructional designers who already use AI in real projects. In eight weeks, in a small group of no more than 15 designers, you build your own AI design process: your context, your step-by-step ways of working (including branching scenarios), your quality standards, and a routine that makes every project better than the last. And you'll see whether it made *you* a better designer, too: you do a short analysis without AI at the start and again in week 8.

2–3 hours a week plus live sessions · max. 15 designers

WHY JOIN THE WAITLIST NOW

Get first pick, before the public launch

The waitlist is free and there's no commitment. You get:

First access before October 20.

The full programme and enrolment details, before the public launch.

The lowest founding prices. The price goes up every five places, so the waitlist has the best chance of a place in the first group:

Places	Founding price	What you get
1–5	€995	Lifetime access + a 1-on-1 session with me after the Lab
6–10	€1,245	Lifetime access
11–15	€1,495	12 months access

All prices excl. VAT. Employer paying? You'll get an invoice and a one-pager for your manager.

Join the waitlist →

<https://markspermon.com/learning-lab/waitlist/>

Mark Spermon: 20 years designing learning for organisations from banks to childcare, 200+ courses, working with AI in my own design process since 2022.