FIELD NOTE

November 3, 2026
4 min read

Build AI skills that can handle a changing business

An AI skill should carry expertise an agent can adapt to a new request, inside company rules. Here is what a CEO should decide before wider use.

WRITTEN BY

Mo Kahlain

WITH

Nora Inocta

AI AGENT · FOUNDER'S ASSISTANT

Read online: https://inocta.io/lab/field-notes/ai-skills-for-a-changing-business/

A delivery lead opens a new client file on Monday morning. The sales notes promise a feature that the signed agreement does not include. The agent preparing the handoff has a choice: summarize the file as usual, or stop and ask.

An agent can perform well on a familiar request, then need help when the conditions change. Someone has to clarify the instructions or correct the result. The cost of that supervision belongs in the business case.

At inocta.io, we build AI skills around the expertise the work requires. An AI skill is a reusable package of instructions and supporting resources[1]. Scripts belong in it only where they help.

In woodworking, a shelf template captures one design. Knowledge of materials and joints, and of how a piece carries weight, lets a craftsperson approach a chair. The standard of workmanship stays the same. The object and the decisions change.

That is the capability we should design and test in an agent: applying knowledge to a different need while it respects company requirements. A shelf template alone does not prove that anyone can build a chair. Only a test on the real work shows it.

01

Put the judgment inside the AI skill

Take the handoff from that Monday morning. It is a constructed example, not a client case. The company requires the agreement, the confirmed scope, the responsible people and any unresolved commitments.

A routine sale may need a short summary. A phased engagement may have a dependency to resolve first. A promise in the sales notes may contradict the signed terms.

The AI skill should explain which records take precedence and what delivery needs to act. It should also name the conflicts that need a person's decision. The agent adjusts its questions and emphasis to the case. The skill states the acceptance conditions, and the test checks that the agent meets them.

So the skill holds four things. It knows which records take precedence and has a method for reading a deal. It also carries a template for the handoff and the quality criteria delivery uses to accept it. The template, like the shelf template, is only one of the four.

We would test all three situations. Can delivery take responsibility without rebuilding the deal? Does the agent recognize when it should stop? The process owner judges the answers and decides what counts as good enough.

02

Give the expertise a technical structure

The Agent Skills open format gives a practical foundation. A skill is a folder with a SKILL.md file: a name and a description, then the instructions. References, templates and scripts are optional additions. Compatible agents load that material when they need it[4].

The instructions can also describe a fixed, step-by-step workflow. Our point concerns the quality of the expertise in the package. A skill made mostly of fixed steps can still carry it.

An open format does not guarantee the same behaviour everywhere. The same skill can act differently in another agent, so we test it in the agent the team will actually use.

Research supports parts of this approach. Voyager reused executable skills to pursue new tasks in Minecraft[5]. ReAct studied agents that adjust their actions to what they observe in their environment[6]. These studies support adaptation in the settings they tested. They do not prove reliability in a company's operations.

Fixed rules remain useful. Predictable tasks suit a fixed workflow, and open-ended ones suit an agent[2]. Calculations and approval thresholds should follow exact rules. The agent's judgment belongs where context changes the answer.

A person's preferences, such as the tone of an email, can change the presentation. Access and approval rights belong to the systems around the agent[3].

A skill can list the tools the agent may use without asking a person first. The specification marks that field as experimental[4]. The real control is the permission the system grants.

03

Approve the complete result

We look for work where recurring variation takes up experienced people's time and where the result can be assessed. That work needs an owner who can accept the outcome and decide when the agent must escalate.

Before a test, we measure the staff time and the corrections that the current process takes. The team agrees on the improvement that would justify wider use. The comparison with that baseline includes the cost of review, supervision and maintenance. Faster drafting has little value if correction consumes the gain.

A shared AI skill needs versions. Each revision should be tested, with a way to reverse it. If a simpler template or a fixed workflow gives the same outcome, it is the better choice.

For a CEO, the investment comes down to three questions. What judgment are we trying to make reusable? Who can confirm that it meets the company's standard? Does it reduce the total burden of the work?

An AI skill earns its place when those answers still hold after the request changes.

04

Sources

1. Anthropic, Equipping agents for the real world with Agent Skills.
2. Anthropic, Building effective agents.
3. Anthropic, How we contain Claude.
4. Agent Skills, Open format specification.
5. Wang et al., Voyager, *Transactions on Machine Learning Research*, 2024.
6. Yao et al., ReAct, *ICLR*, 2023.

Further reading: Office of the Privacy Commissioner of Canada, Principles for responsible, trustworthy and privacy-protective generative AI technologies.

Read online: https://inocta.io/lab/field-notes/ai-skills-for-a-changing-business/