

ANATOMY OF AN AI SKILL

Example: [ai-haavoittuvuus.skill](#)

How a long prompt becomes a portable, evidence-backed and repeatable operating model.

SKILL.md



ASSETS



REFERENCES



EXAMPLES

12 files. 7 stages. 1 repeatable outcome.

This minibook opens the package and explains it layer by layer.

Not one prompt. A 12-file package.

The ZIP carries both the instructions and the implementation memory.

ai-haavoittuvuus.skill

```
ai-haavoittuvuus/  
  
|- SKILL.md  
  
|- assets/  
  
|   |- itsearviointi-pohja.md  
  
|   |- video-ajolista.md  
  
|   |- gmb_logo_transparent.png  
  
|   '- mape_photo.png  
  
'- references/  
  
    |- arviointimalli.md  
  
    |- exposure-baseline.json  
  
    |- anthropic-economic-index.md  
  
    |- mstc-ja-life-skills.md  
  
    |- toimenpidekehys.md  
  
    |- headai-demo.md  
  
    '- esimerkki-markku.md
```

1

CONTROL LAYER

SKILL.md

4

BRAND + USE

assets

7

DATA + METHOD

references

Core idea

The skill can explain itself, locate its sources and show what good looks like.

SKILL.md defines when and how to use the skill.

Other files reveal deeper knowledge only when it is needed.

name: ai-haavoittuvuus

description:

AI vulnerability diagnosis - assess a person's or role's exposure task by task and turn the vulnerability map into an automation map.

1

TRIGGER

AI vulnerability, exposure analysis, tasks likely to automate...

2

OUTCOME

Matrix, indices, baseline, actions and a GMB-branded diagnosis.

3

PROCESS

Seven stages from task inventory to visual QA.

4

BOUNDARIES

No fearmongering. No invented data. Every risk gets a mitigation.

One skill. Seven guided stages.

Each stage says what to do and where to find the supporting detail.

1

TASK INVENTORY

8-12 tasks + time allocation

2

SCORING

A and H axes, four quadrants

3

BASELINE

theoretical vs. observed exposure

4

DATA DEMO

HeadAI demand signal, optional

5

THE TURN

from threat to automation portfolio

6

MSTC + LIFE SKILLS

development path and gold skills

7


OUTPUT + QA

GMB document, rendering and review

Skill quality is not created by instructions alone. It comes from the chain of instructions, sources, examples and QA.

Decompose the role into tasks.

Vulnerability lives in the weekly task mix, not in the job title.



USEFUL TASK

'Training material production'

Concrete, repeatable and possible to estimate in time.


TOO BROAD

'Client work'

Hides tasks with very different exposure profiles.

Two axes. Four action choices.

A = automation potential. H = human premium.



The goal is not to escape red. The goal is to automate it yourself first and reinvest the released time in gold.

Three metrics turn analysis into a decision.

Example values come from the skill's internal Markku demo, not general statistics.

52 /100

VULNERABILITY

time-weighted automation potential

65 /100

PROTECTION

share of the week protected by human premium

~8 h/week

THE TURN

time potentially released through automation


A baseline does not give a verdict. It gives a window to act.

- **THEORETICAL** What AI could theoretically do
- **OBSERVED** What is already being done with AI
- **SELF-ESTIMATE** What moves now and within 12 months

Theoretical - observed gap = time to act, not a forecast of job loss.

Read the vulnerability map a second time.

Red tasks are not only risks. They are your first automation portfolio.




- 0-3 MO** Automate 1-2 red tasks and use AI to augment one yellow task.
- 3-12 MO** Build a personal AI-OS, strengthen gold skills and measure the time shift.
- 12-24 MO** Reposition the role around human premium and scale automations to the team.

Risk without mitigation is anxiety, not analysis. Every risk observation gets a next action in the same output.

Why does the skill move from one model to another?

Because the expertise lives in files, not in one chat history.



OpenAI describes a skill as a directory with SKILL.md plus optional scripts and references. The structure builds on the open Agent Skills standard.

YOUR NEXT SKILL

STORE YOUR EXPERTISE. NOT JUST THE CHAT.



How cool is that?

- 1. Write a clear trigger and outcome.
- 2. Move sources and examples into folders.
- 3. Make the process staged and verifiable.
- 4. Package it, run it in two environments and compare.

Markku Pelkonen
GrowthMindsetBuilder

markku.pelkonen@growthmindsetbuilder.com
+358 45 318 7300

Sources: ai-haavoittuvuus.skill; Anthropic Economic Index report: Cadences, 26 Jun 2026;
OpenAI Codex, Build skills. Re-check external figures before any client diagnosis.