SKILL PUBLISHER

What a skill costs to carry

Every skill your agents hold is paid for on each request that pulls it in.

The problem it solves

A skill is not free to own. It occupies the same context window as the work, and it loads on whatever its description claims rather than on whatever it can actually help with. A library assembled one skill at a time ends up with descriptions written to the character limit and one ordinary question pulling four of them.

Measured	Cut, not trimmed	Nothing operative lost
tokens, triggers, overlap	whatever did not earn its place	every deletion read first

Where the weight comes from

1. **Descriptions written to the limit**
   Each author maximises their own skill's trigger coverage in isolation, which is the sensible thing to do alone. The result is a family where one general question pulls four.
2. **Wrappers that restate the obvious**
   A section on when to use the skill, an essay on why a machine cannot do this work, a changelog with its reasoning attached. All loaded every time, none of it changing an answer.
3. **Triggers that enumerate**
   A description listing the tool names that appear in the question gets pulled in by the list, including for every question it cannot serve.
4. **Restated general knowledge**
   The part a capable model produces unprompted. It is the largest category, it is invisible without a with-and-without test, and it is why measuring comes before editing.

What the work returns

- A library that costs less to carry, with the rules that mattered still in it.
- The evidence for every cut kept beside the skill, rather than in somebody's memory.
- Descriptions that say what a skill is not for, so it stops loading defensively.
- A repeatable pass, so none of this has to be discovered again next year.

Our own library, before and after

- 265,915 tokens if every skill loaded at once, before the pass. 212,982 after, with no rule lost.
- 48,399 tokens recovered across 160 skills, and about forty operative lines rescued that existed nowhere else and would have gone with the wrapper.
- 62 descriptions were over the format's hard limit, where a truncation takes the exclusions first, because they are written last.

Measured September 2026 on this library. Context for what this work recovers, not a claim about what yours will.

We assess skills you already run, whoever wrote them.

skillpublisher.com

To find out more: sales@skillpublisher.com