

pxpipe: Turning AI Agent Prompts Into Images to Cut Token Costs

Kabui, Charles

2026-07-20

[Read at ToKnow.ai](#)



[pxpipe 0.9.0](#) is a local proxy that turns bulky AI agent input into PNG images before sending it to a model. System instructions, tool descriptions, old chat history, and dense tool results can all take the image route, while replies remain text. The project's current example renders 48,000 characters as about 2,700 image tokens instead of 25,000 text tokens. Its logs show [59% to 70% lower total bills](#), depending on the workload. In one maintainer-run comparison,

the imaged session completed the same counting and arithmetic task for [\\$4.51 instead of \\$42.21](#), although it needed one extra prompt to follow the requested answer format. These are project-run results, not an independent benchmark.

Coding agents repeatedly resend long instructions, tool definitions, files, and earlier messages. Converting that bulk can cut the bill without changing the model or shortening its reply, especially when the client has not already cached most of the prompt. The catch is accuracy. A [blind reading audit](#) recovered only 15 of 24 exact strings, or 63%, from dense images. pxpipe keeps recent turns and selected identifiers as text, and unsupported models pass through unchanged. Because it is a proxy, it [handles the full request](#). Optional PNG dumps can also preserve sensitive context. This works best for repetitive, text-heavy sessions, while exact code, hashes, and secrets still need a text path.

Read More: [five other ways to cut AI agent token costs](#).

Sources:

- [pxpipe 0.9.0 on npm](#)
- [pxpipe repository and benchmarks](#)
- [Effective-context comparison attempt log](#)
- [Imaged-text legibility audit](#)
- [Prompt caching and savings accounting](#)

Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)