

A Typst HTML Paper Pipeline for Astro

A small architectural note

hanbings
Independent Researcher
hanbings@hanbings.io

August 4, 2026

Abstract. This note demonstrates a dual-target content pipeline in which Typst emits semantic HTML for reading inside an Astro blog and a paged PDF for archival and print use. The renderer remains distinct from Markdown without becoming visually detached from the surrounding website.

Keywords: Typst; Astro; HTML; academic publishing; static sites.

1. Introduction

Academic writing and personal publishing optimize for different reading modes, but they do not need to live on visually disconnected websites. A paper can use the blog’s navigation, colors, and reading width while retaining its own semantic structure for equations, figures, references, and academic metadata.

Typst is a programmable typesetting system designed for structured documents and mathematical notation. Its paged output can preserve expressions such as

$$\int_0^\infty e^{-x^2} dx = \frac{\sqrt{\pi}}{2}$$

as MathML in the browser and as precisely typeset mathematics in PDF. Astro can therefore remain responsible for the surrounding reading experience without turning the paper into Markdown.

1.1. Separation of responsibilities

The source tree assigns one job to each layer:

- `paper.typ` owns the paper body and its local project files.
- `paper.json` owns web metadata such as the abstract and canonical date.
- Typst emits semantic HTML for the website and a PDF for download.
- Astro injects the HTML body into the normal blog shell and styles it separately.

This separation follows a familiar principle: source formats should retain their native semantics instead of being coerced into a shared lowest common denominator.

2. Build model

Let S denote the set of files in one paper project and let V denote the Typst compiler version. The build artifact can be modeled as

$$A = \text{compile}(\text{hash}(S), V).$$

The content digest makes unchanged papers reusable across repeated Astro checks and builds. A changed bibliography, image, chapter, or main source invalidates the same paper without rebuilding unrelated entries.

Layer	Responsibility
Typst	Semantic HTML, paged PDF, equations, figures, bibliography
Astro	Metadata, URLs, navigation, responsive publication shell
Browser	Native text, links, MathML, and downloadable artifacts

Table 1: Responsibilities of the paper pipeline.

3. Discussion

Typst currently emits a standalone HTML document rather than an embeddable fragment. The build loader extracts its body and compiler-provided styles, then Astro places that semantic output inside the site’s own layout. Readers keep the normal blog navigation and can select text, follow links, and inspect equations.

HTML export is still experimental, so the PDF remains the stable archival artifact. Target-aware rules keep print-only title geometry and page breaks out of the flowing web version while allowing both outputs to share the same paper body.

4. Conclusion

A paper-style post renderer is a content contract, not a separate visual world. Markdown remains optimized for essays, while Typst becomes a first-class source format for academic work inside the same site.