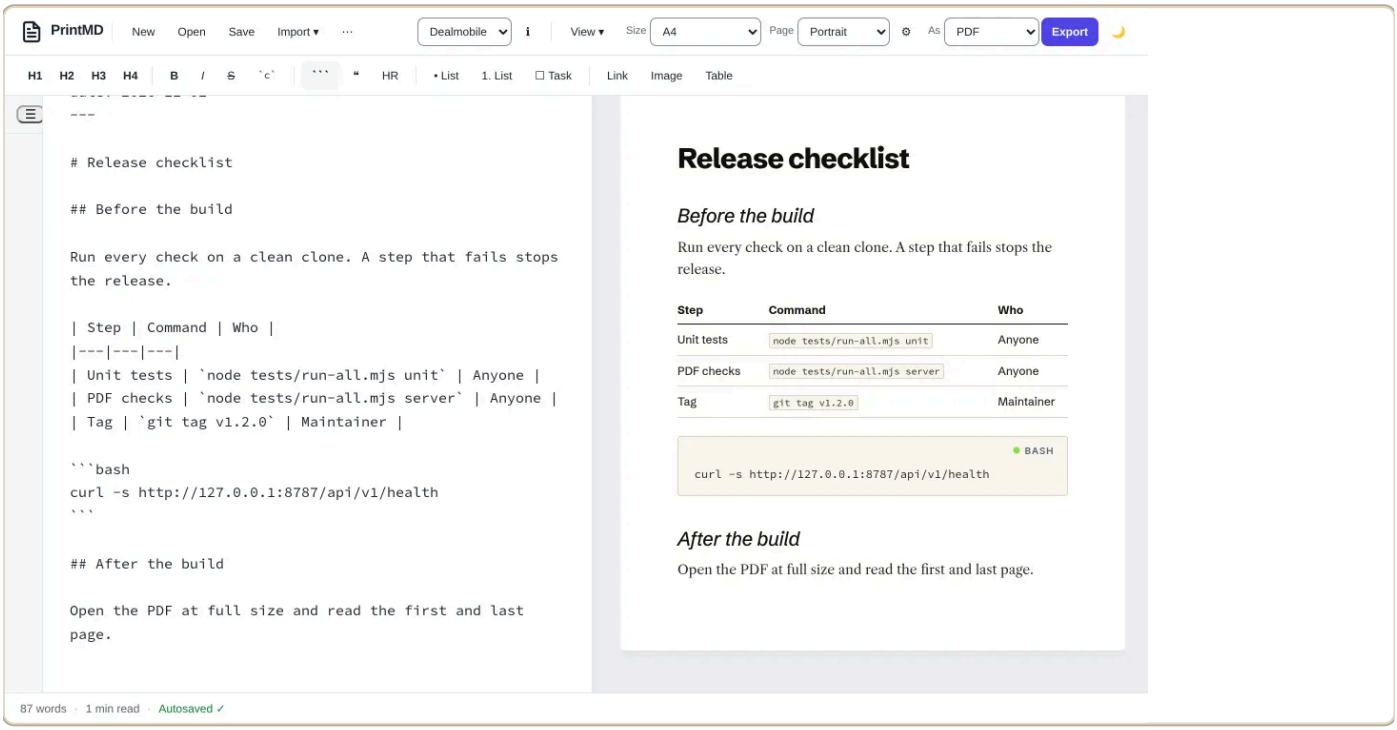
Print-quality PDFs from a browser

A browser can print any page to PDF, but a document you send someone needs more. How PrintMD gets fonts, page breaks and images right, and what its tests check.



Published
2 November 2026

Reading time
3 minutes

About
PrintMD

Tags
Markdown, Testing

More on: PrintMD

A browser can make a PDF from any page, and for a quick copy of a web page that's fine. For a document you send to someone, "fine" isn't the bar. This post is about the gap between the two, and what PrintMD does to close it.

Start with the fonts

If a PDF names a font the reader doesn't have, their viewer swaps in another one and your layout moves. The cure is to embed the font in the file, and for that it has to be a real file the page can load, not a name in a stylesheet.

PrintMD bundles its fonts, all under the SIL Open Font Licence, and loads them from the app itself. Before printing, every bundled face the text uses is loaded and checked. If one didn't load, you're told which one, rather than getting a system font by surprise. That covers Latin text: the bundled fonts don't include Cyrillic, Greek or Vietnamese, so those, and emoji, still come from the machine's own fonts, without a warning.

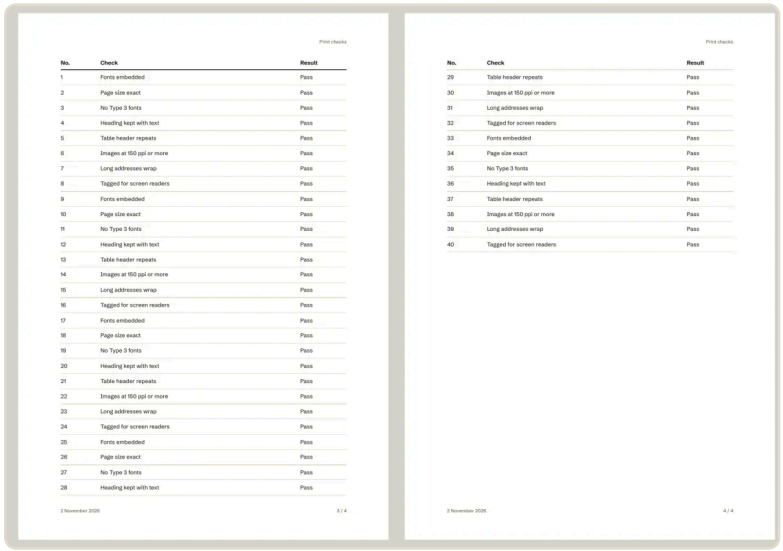
Bold and italic need real faces too. When a family has no bold, the browser makes one up, and a made-up bold prints as a Type 3 font, which is the kind of font the tests reject. Some of the bundled families have a single italic, so their bold italic prints as that italic at the regular weight rather than as a true bold italic.

Page breaks are the other half

Most of what makes a PDF look careless is where the pages break. A few rules cover most of it. They're requests the browser keeps wherever it can, so a table row taller than a whole page still has to break:

- A heading isn't left at the foot of a page.
- At least three lines of a paragraph stay on either side of a page break.
- A table row, a figure and a short code block stay whole.
- A table that runs onto the next page repeats its header row.
- An image is never taller than the text column.
- Long web addresses and hashes wrap instead of running off the page.

None of these is clever. They're CSS print rules, and the point is to apply them in every template, every time.



No.	Check	Result
1	Fonts embedded	Pass
2	Page size exact	Pass
3	No Type 3 fonts	Pass
4	Heading kept with text	Pass
5	Table header repeats	Pass
6	Images at 150 ppi or more	Pass
7	Long addresses wrap	Pass
8	Tagged for screen readers	Pass
9	Fonts embedded	Pass
10	Page size exact	Pass
11	No Type 3 fonts	Pass
12	Heading kept with text	Pass
13	Table header repeats	Pass
14	Images at 150 ppi or more	Pass
15	Long addresses wrap	Pass
16	Tagged for screen readers	Pass
17	Fonts embedded	Pass
18	Page size exact	Pass
19	No Type 3 fonts	Pass
20	Heading kept with text	Pass
21	Table header repeats	Pass
22	Images at 150 ppi or more	Pass
23	Long addresses wrap	Pass
24	Tagged for screen readers	Pass
25	Fonts embedded	Pass
26	Page size exact	Pass
27	No Type 3 fonts	Pass
28	Heading kept with text	Pass

No.	Check	Result
29	Table header repeats	Pass
30	Images at 150 ppi or more	Pass
31	Long addresses wrap	Pass
32	Tagged for screen readers	Pass
33	Fonts embedded	Pass
34	Page size exact	Pass
35	No Type 3 fonts	Pass
36	Heading kept with text	Pass
37	Table header repeats	Pass
38	Images at 150 ppi or more	Pass
39	Long addresses wrap	Pass
40	Tagged for screen readers	Pass

A real render: rows 29 to 40 carry on at the top of page 4 under the repeated header.

The tests guard some of these better than others. Wrapping long lines, addresses and hashes is checked in a real PDF for every template. The heading, table header and tall image rules are checked in a real PDF in the Unbranded template only. The three-line rule and the table row and figure rules are in the stylesheet, and no test prints them yet.

Effects that turn into pictures

A soft shadow surprised me. The browser draws a blurred shadow as a picture, so a decorative card with a soft shadow becomes images inside the PDF, and the file grows. Blur filters do the same. A shadow with no blur, or a see-through fill, stays as vector shapes. So the templates avoid blurs and filters, and the tests count the images in each sample PDF and fail if there are more than the document holds.

Images are left as they are

PrintMD doesn't resize or redraw images: a 6,000-pixel-wide PNG goes into the PDF at 6,000 pixels. It does work out the resolution each one will print at, and lists any under 150 pixels per inch, because that's where print starts to look soft. SVG is vector, so it isn't checked.

Checking it

The tests print sample documents through the PDF service, in every template, and look at the result with the usual poppler tools: the page size, whether every font is embedded and none is Type 3, the images, the tags for screen readers, the title, the language and the bookmarks. That catches the regressions a glance at the preview wouldn't. What they don't check is reading order: I haven't written a test for it.

To see it for yourself, run the editor, open a long document with a table in it, and export it as a PDF.