Vocalised right-to-left text

Nikud and harakat are combining marks. They have no width and no place of their own: the face states an anchor on the letter and an anchor on the mark, and the two are brought together by GPOS. tclpdf reads those anchors, so the pointed spelling is set instead of refused - and because a mark carries no advance, it is set in exactly the same room as the bare one.

Hebrew

with nikud	7 characters, 18.736 mm	שָׁלוֹם
without	4 characters, 18.736 mm	שלום

Seven characters against four, and the same width to the thousandth of a millimetre. The points are drawn where the anchors of DejaVu Sans put them - measured against HarfBuzz, 366 of 366 consonant-and-point pairs of that face agree to the unit. Not every face is that clean: in Liberation Sans, Liberation Serif and Arimo, 25 of 1188 pairs sit 25 units of 2048 out - a hundredth of an em, all of them the holam, because those faces refine it in a chaining contextual lookup (GPOS type 8) that tclpdf does not read.

Arabic, in two faces that are built differently

DejaVu Sans - whole letters	9 characters, 19.575 mm	مَرْحَبًا
	5 characters, 19.575 mm	مرحبا
Noto Naskh Arabic - skeleton plus dot glyphs	9 characters, 17.280 mm	مَرْحَبًا
	5 characters, 17.280 mm	مرحبا

Both faces set the word with its harakat, and the second one is the interesting half: every dot under the beh and over the noon of that face is a glyph of its own, made by ccmp and placed by the same anchors. A face of that kind was refused while mark attachment was missing, because its dots would have landed beside their letters rather than on them.

Tibetan is still refused, for a different reason

ཀཁགངཅ

tclpdf: U+0F90 (position 1) belongs to Tibetan, which needs stacked forms out of GSUB - tclpdf does not do that. Set -unshaped 1 to draw it as isolated glyphs in logical order anyway

the same syllable with -unshaped 1: ས྄ག

A single vowel sign on a single letter is placed here exactly as HarfBuzz places it - 160 of 160 pairs of Noto Serif Tibetan. It is the syllable that fails: HarfBuzz replaces letter and mark with one precomposed stack glyph, so "skad" is two glyphs there and three here. Anchors cannot supply a glyph the character map does not lead to, which is why the whole block of stacking marks stays refused and the message names GSUB rather than mark placement.

Check it yourself

```
qpdf --check 02.12-vocalised-rtl.pdf
qpdf --qdf --object-streams=disable 02.12-vocalised-rtl.pdf - | grep TJ
pdftotext 02.12-vocalised-rtl.pdf - | head -20
```