

When compliant isn't accessible

Why passing PDF/UA validation doesn't guarantee a usable document — and how AccessiblePDF closes the gap.

A Palanisi Labs white paper · accessiblepdf.app · June 2026

Executive summary

Organizations are under a hard, expanding legal mandate to make their PDF documents accessible — yet most remediation tools optimize for a single thing: passing an automated validator. A document can pass PAC or veraPDF and still be unusable with a screen reader. Across more than 400 production documents, we found this gap to be the rule rather than the exception, and it concentrates in the document types that matter most: government tables, large legal texts, and scanned files.

AccessiblePDF is built around the outcome that actually matters — what a screen-reader user experiences — and validated against it. In a controlled test of nine previously unseen documents spanning 189 pages, it reached an average PAC score of 98.1 and remediated 1,275 of 1,276 detected accessibility issues, with eight of nine documents reaching a perfect score.

A document can pass every automated check and still be unusable with a screen reader. Compliance is necessary. It is not sufficient.

A mandate that is no longer optional

In April 2024 the U.S. Department of Justice issued a final rule under Title II of the Americans with Disabilities Act requiring state and local governments — including public universities — to bring their websites, applications, and digital documents up to the WCAG 2.1 Level AA standard. After a one-year extension in 2026, the compliance deadlines now fall on April 26, 2027 for larger public entities and April 26, 2028 for smaller ones. Section 508 covers federal agencies, the European Accessibility Act covers the EU, and sector rules add still more. PDFs and other documents are explicitly within scope.

The scale of the problem is enormous: there are trillions of PDFs in circulation, and an estimated nine in ten contain accessibility barriers. The obligation does not stop at a government's own files. The contractors, vendors, nonprofits, and institutions that submit documents to public agencies must also deliver accessible PDFs — and some agencies have already begun rejecting submissions that do not comply, ahead of formal deadlines. Where the required deliverable is a PDF, converting to accessible HTML is not an accepted substitute.

Tellingly, when the DOJ extended its deadline, it cited the limits of current technology — including generative AI — to automate accessibility remediation at scale. The regulator itself has acknowledged that the problem is not yet solved.

The hidden gap: passing the checker is not the same as usable

Automated validators such as PAC and veraPDF answer a specific question: does the document contain the required accessibility structure, and is it well-formed? They confirm that tags are present, that a structure tree exists, that the language is declared, that roles are mapped. This is necessary work, and these tools are indispensable. But they largely cannot answer the question that matters to a person: is the document actually navigable with assistive technology?

A file can satisfy every automated checkpoint and still fail a real screen-reader user. In practice, the failures cluster in predictable places:

- Reading order that does not follow the document's logical flow.
- Tables that are tagged as tables but whose cell and header relationships are broken, leaving a screen-reader user unable to navigate them.
- Heading structures that are present but incorrect.
- Text that extracts as gibberish because of broken font encoding — the validator sees that text exists, while the screen reader announces nonsense.
- Scanned, image-only pages with no recoverable structure at all.

These problems are not evenly distributed. They concentrate in the hardest, highest-stakes documents — government procurement and bid packages dense with financial and compliance tables, lengthy legal texts, and scanned contracts. Those are precisely the documents public agencies receive every day.

A different design principle

Most remediation follows a validator-first loop: detect, tag, pass the checker, ship the file. AccessiblePDF is built on a different principle — it optimizes for the reader, not the checker. Each document is remediated for correct reading order, genuine table semantics, accurate heading hierarchy, and meaningful descriptions of images, while preserving the document's exact original appearance rather than rebuilding it.

Every remediated file is then validated two ways: against the PAC and veraPDF reference standards, and through manual review in the NVDA screen reader. Validator-pass and human usability are different tests, and we run both.

Evidence

This approach did not come from theory. It emerged from analyzing more than 400 production documents, which is where the divergence between automated compliance and real usability became clear. To validate first-pass performance on documents the engine had never seen, we ran a controlled test of nine fresh documents totaling 189 pages — including a 70-page government procurement package and a 47-page legal text.

| Document type | Pages | Issues remediated | PAC |
|--------------------------------------|-------|-------------------|-----|
| Technical e-book | 25 | 303 / 303 | 100 |
| Government bid / procurement package | 70 | 213 / 213 | 100 |
| Product datasheet | 6 | 205 / 205 | 100 |

| Document type | Pages | Issues remediated | PAC |
|-------------------------------------|-------|-------------------|-----|
| Professional résumé | 2 | 204 / 204 | 100 |
| Training booklet | 8 | 188 / 188 | 100 |
| State constitution / legal document | 47 | 108 / 108 | 100 |
| Government RFP template | 12 | 38 / 38 | 100 |
| Construction contract | 16 | 11 / 11 | 100 |
| Reference list / table | 3 | 5 / 6 | 83 |

One additional document was excluded from the set because its source file was corrupt and would not open — we report that rather than quietly dropping it. The single document that did not reach a perfect score was a short reference table where one of six detected issues remained. The realism of these numbers is the point: a tool that claims a flawless result on every document invites the skepticism it deserves.

We hold our own work to the same standard. This white paper, and the validation report behind it, were themselves remediated by AccessiblePDF.

Why it matters now

For an organization that submits documents to a public agency, an inaccessible PDF is no longer a quality issue — it is a rejected submission against a fixed deadline. AccessiblePDF turns what has been a slow, specialist, and expensive task into a fast, self-serve one. Professional manual remediation commonly runs from several dollars to twenty-five dollars per page; AccessiblePDF delivers compliant, screen-reader-usable output at a fraction of that cost and without specialist labor.

For the agencies on the receiving end, the same capability offers a way to point submitters toward a tool that produces genuinely usable documents — and to reduce the volume of non-compliant submissions they must turn away.

Conclusion

Compliance is the floor. Usability is the point. The standards exist to serve people who rely on assistive technology, and a document that passes every automated check but cannot be read by a screen-reader user has met the letter of the requirement while missing its purpose. AccessiblePDF is built for the purpose.

AccessiblePDF is a product of Palanisi Labs. Learn more at accessiblepdf.app.