

WHITE PAPER · DOCUAGENT ARCHITECTURE SERIES

WCAG 2.2 Accessibility Auditing at Enterprise Scale

How continuous, automated auditing replaces the six-figure manual accessibility retainer

Document ID: DA-WP-2026-004

Cryptographic boundary: FIPS 140-3 validated

DocuAgent AI

Prepared for compliance, security, and engineering leadership

Executive Summary

Enterprise SaaS companies in regulated industries — finance, health tech, government, and digital commerce — are running into a wave of overlapping accessibility rules at once: Section 508 requirements in US federal procurement, ADA Title III enforcement from the Department of Justice, and the now-binding European Accessibility Act under EN 301 549. WCAG 2.2 Level AA compliance isn't something a team checks off once a year anymore. It has to hold up every time code ships.

For years, the standard approach was to bring in an outside accessibility consultancy. Those audits typically run \$40,000 to \$120,000 per application, take six to twelve weeks to complete, and produce a static PDF report that's already out of date by the time the next release goes out.

This paper covers the architecture behind DocuAgent AI's approach to the problem: a microVM-isolated, vision-augmented, headless system that audits WCAG 2.2 AA continuously, inside the same pipelines engineering teams already use. It runs inside a FIPS 140-3 compliant cryptographic boundary and replaces the old audit-retainer model with real-time accessibility scoring, visual step-by-step evidence, and gating that catches regressions before they reach production.

The Regulatory and Operational Challenge

What Non-Compliance Actually Costs

Since the European Accessibility Act took effect across EU member states, companies with non-compliant digital products are exposed to fines of up to €100,000, or 4% of global annual turnover, along with the possibility of being pulled from the market entirely. In the US, ADA Title III lawsuits over digital accessibility and Section 508 procurement rejections can quietly stall an entire public-sector sales pipeline for a vendor that can't demonstrate compliance.

Why the Old Audit Model Doesn't Hold Up

Manual accessibility auditing runs into the same three walls once it hits enterprise scale:

- **It's expensive and slow.** Manual consultancies test with actual screen reader users working through the application by hand. A 50-route web application can take more than 150 human hours to audit, which rules out testing every deployment.
- **The results go stale fast.** Teams shipping code daily can make a manual audit deliverable outdated within hours of it being published.
- **Automated scanners barely scratch the surface.** Off-the-shelf accessibility linters catch roughly 30% of WCAG failures. They can read static DOM syntax, but they miss dynamic single-page app state changes, focus traps that only appear during interaction, and contrast issues on canvas or background elements.

Technical Architecture: How It Works

DocuAgent AI builds WCAG 2.2 evaluation directly into a headless, continuous monitoring pipeline, rather than treating it as a separate exercise someone runs before an audit deadline. A scan moves through four layers:

Stage	Pipeline Layer	What Happens
1	Target Enterprise SaaS Application	The system points at the live application exactly as an end user would experience it — no staging shortcuts.
2	FIPS 140-3 Boundary / gVisor MicroVM Worker	AST and DOM parsing, a vision-based canvas contrast evaluator, a keyboard navigation state machine, and client-side PII redaction all run inside one isolated worker.
3	DocuAgent Audit Engine	WCAG 2.2 AA violations get scored, SVG step-pins mark exactly where each issue occurs, and the results are written to a SHA-256 tamper-evident vault.
4	Continuous Output & CI/CD Gating	Visual diff comments post directly to GitHub pull requests, and Section 508, VPAT, and EAA evidence bundles are generated automatically.

WCAG 2.2 Coverage: Where Automation Replaces Manual Testing

DocuAgent pairs DOM abstract syntax tree (AST) analysis with computer vision to handle WCAG 2.2 Level A and AA criteria that used to require a person sitting with the application. A few examples:

Success Criterion	Level	Why It's Hard to Audit by Hand	How DocuAgent Automates It
2.4.11 Focus Not Obscured (Minimum)	AA	Sticky headers or footers can cover the element that currently has keyboard focus.	The microVM worker compares the bounding box of sticky elements against the focused element's bounding box to catch overlap.
2.4.13 Focus Appearance	AA	Confirming a focus indicator meets the 3:1 contrast ratio and 2px minimum thickness takes pixel-level comparison.	The vision engine compares the rendered canvas before and after focus to measure the contrast delta and outline thickness.
2.5.8 Target Size (Minimum)	AA	Touch targets need to stay at or above 24x24 CSS pixels across every responsive breakpoint.	AST traversal computes the bounding rectangle of every interactive element and flags anything under 24px.
3.3.7 Redundant Entry	A	Multi-step forms sometimes force people to re-enter information they already submitted.	A graph-based navigation engine tracks each field's signature across steps and flags duplicate requests for the same data.
3.3.8 Accessible Authentication	AA	Login flows can unintentionally block password managers or require solving a memory-based puzzle.	The DOM inspector checks for autocomplete attributes, copy-paste support, and WebAuthn or passkey compatibility.

FIPS 140-3 Compliance and Zero-Trust Execution

Auditing regulated healthcare (HIPAA) and financial (GLBA / SOC 2) platforms means the audit tooling itself has to meet a high bar. DocuAgent runs under a few specific controls to make that possible:

- **FIPS 140-3 validated modules.** Every piece of evidence data, at rest and in transit, is encrypted using FIPS 140-3 Level 1/2 validated algorithms — AES-256-GCM for the data itself and SHA-256 hash chains for the audit log.
- **gVisor and Firecracker microVM isolation.** Crawling workers run inside disposable microVM sandboxes. No session data or auth tokens survive past the end of a run.
- **PII redaction before anything leaves the sandbox.** Screenshots and DOM nodes are scrubbed of Social Security numbers, medical record numbers, and credit card numbers by client-side regex and AST rules before they ever reach a vision model or the audit vault.

Enterprise ROI and Deployment Outcomes

Moving from a periodic manual consultancy to DocuAgent's continuous model tends to show up in three places:

- **Audit spend drops by around 85%.** Continuous automated coverage replaces the recurring six-figure retainers most enterprises pay to outside consultancies each year.
- **Regressions get caught before they merge.** A pull request fails automatically the moment a change drops a focus indicator or shrinks a touch target below 24px.
- **VPAT and EAA evidence is ready on demand.** Every scan produces a signed, SHA-256-backed package with visual step pins, DOM element paths, and remediation notes — ready to hand to a procurement officer or an outside auditor without extra prep work.

Where DocuAgent Fits Into Your Stack

DocuAgent is meant to sit inside the tools an engineering team already has open, rather than becoming one more dashboard someone has to remember to check. A few of the integration points that come up most often:

- **MCP server support.** Connect DocuAgent to any MCP-compatible client, so it can be called directly from an assistant or an internal tool instead of running as a standalone app.
- **GitHub Action PR diff comments.** Every scan can post its findings straight into the pull request that triggered it, so a reviewer sees a broken focus indicator or an undersized touch target right next to the code that caused it.
- **Splunk and Datadog webhooks.** Violation alerts and evidence events stream into whichever SIEM or observability platform is already in place, instead of living in a report only the accessibility team ever opens.
- **Local vLLM/Ollama fallback.** When a scan needs model inference and sending page content to an external API isn't appropriate, DocuAgent can fall back to a locally hosted vLLM or Ollama instance instead.

See a Sample Evidence Bundle

The easiest way to get a sense of what a DocuAgent audit actually looks like is to open one. A sanitized sample Compliance Evidence Bundle is available to browse directly on docuagent-ai.com — no account or sales call required. It's a real .zip export, built against a demo application, with the same step-pins, scoring, and signed ledger structure a production scan would produce.

Inspect a sample bundle at docuagent-ai.com, or browse the rest of this series and full integration guides at docuagent-ai.com/resources/whitepapers and docs.docuagent-ai.com.