OB.1 AI SOLUTIONS

ADA Website Compliance Scanner

Technical Research Brief

Legal framework, technical standards, detection tools, agent architecture, violation taxonomy, and market landscape for building an AI-native accessibility scanning and remediation agent.

Prepared by OB.1 AI Solutions

chris@ob1ai.co · ob1ai.co · 234.602.0500

April 2026

From Assessment to Architecture.

TABLE OF CONTENTS

1. Executive Summary

5,114	57%	95.9%	$1M
ADA lawsuits filed (2025)	Issues caught by automation	Top 1M sites failing checks	FTC overlay fine (2025)

The DOJ's April 2024 Title II rule formally adopted **WCAG 2.1 AA** as the first binding web accessibility standard under the ADA, but only for government entities with a compliance deadline of **April 24, 2026**. For private businesses under Title III, no formal technical standard exists. Yet 5,114 ADA digital accessibility lawsuits were filed in 2025 (a 20%+ increase from 2024), with plaintiff firms using automated tools to identify targets at industrial scale.

Automated testing catches roughly 57% of real-world accessibility issues by volume (Deque study of 300,000+ issues), though only ~30% of WCAG success criteria are fully automatable. The overlay/widget industry has been definitively discredited: the FTC fined AccessiBe $1 million in April 2025 for false advertising, and 25% of all 2024 ADA lawsuits targeted websites already running overlay tools.

This creates a significant market gap: SMBs need affordable, real remediation between useless $490/year overlays and $100-500/page manual audits. An AI-native scanning agent built on Playwright + axe-core, augmented with LLM-powered judgment for the 60%+ of criteria requiring human review, could fill that gap.

The Enforcement Paradox

The Trump administration's January 2025 deprioritization of ADA enforcement created a paradox: reduced federal enforcement has accelerated private litigation. The DOJ rescinded 11 ADA guidance documents on March 19, 2025, and announced in September 2025 it would not pursue pending ADA rulemakings. Yet federal website-specific ADA lawsuits surged to 3,117 in 2025 (27% increase), while total digital accessibility lawsuits reached 5,114. Pro se plaintiffs using AI tools to draft complaints increased 40%.

The top 10 plaintiff law firms filed over 80% of federal cases. Multiple legal authorities (Ogletree Deakins, the ABA, Seyfarth Shaw) confirm the Title II rule creates strong precedent pressure for Title III private-sector application. Courts already reference WCAG 2.1 AA as the benchmark in Title III cases despite no formal codification.

The Safe Harbor Question

> **There is no formal legal safe harbor.** Even the Title II rule explicitly states that WCAG conformance does not immunize entities from all claims. However, WCAG 2.1 AA compliance is the strongest defensible position available. Courts routinely order WCAG compliance as an equitable remedy (as in Robles v. Domino's), and meeting the standard substantially reduces legal exposure. The gap between "no safe harbor exists" and "this is what every court references" represents a core legal gray area any AI agent should help clients navigate.

Automation Ceiling

Automated tools detect ~30% of WCAG success criteria fully and ~44% partially, leaving ~36% requiring human review. However, measured by issue volume, automation catches 57% of real-world problems because the most

common violations (missing alt text on 55.5% of sites, low contrast on 79.1%, missing form labels on 48.2%) happen to be automatable and occur at enormous scale. The WebAIM Million 2026 study found an average of 56.1 errors per homepage, with 95.9% of the top million sites failing automated checks. An AI layer analyzing axe-core's "needs review" items could push practical coverage toward 70-80%.

2. Standards Reference Sheet

WCAG 2.1 Level A + AA Success Criteria (50 total)

The following table documents all 50 WCAG 2.1 Level A and AA success criteria with automatability ratings, detection methods, and priority levels for the scanning agent. Bolded IDs represent the highest-lawsuit-risk items.

WCAG ID	Name	Level	Auto?	Detection Method	Priority
1.1.1	Non-text Content	A	Partial	Missing alt: auto; quality: AI/manual	Critical
1.2.1	Audio/Video-only	A	No	Manual transcript verification	High
1.2.2	Captions (Prerecorded)	A	Partial	Auto: video without track; quality: manual	Critical
1.2.3	Audio Description	A	No	Manual verification	High
1.2.4	Captions (Live)	AA	No	Real-time human evaluation	Medium
1.2.5	Audio Desc. (Prerecord)	AA	No	Manual verification	Medium
1.3.1	Info and Relationships	A	Partial	Auto: labels, headings, tables; semantics: manual	Critical
1.3.2	Meaningful Sequence	A	Partial	DOM vs. visual order comparison	High
1.3.3	Sensory Characteristics	A	No	Manual content review	Medium
1.3.4	Orientation	AA	Full	CSS orientation-lock detection	Low
1.3.5	Identify Input Purpose	AA	Full	Autocomplete attribute validation	Medium
1.4.1	Use of Color	A	Partial	Link contrast check; color-only: manual	High
1.4.2	Audio Control	A	Full	Autoplay audio element detection	Medium
1.4.3	Contrast (Minimum)	AA	Full	Luminance ratio calc (4.5:1 / 3:1)	Critical
1.4.4	Resize Text	AA	Partial	Viewport meta check; zoom: manual	Medium
1.4.5	Images of Text	AA	Partial	OCR detection; many false negatives	Medium
1.4.10	Reflow	AA	Partial	Responsive behavior; content loss: manual	Medium
1.4.11	Non-text Contrast	AA	Partial	UI component contrast measurement	High
1.4.12	Text Spacing	AA	Full	Inline spacing override detection	Medium
1.4.13	Content on Hover/Focus	AA	No	Dismissibility/persistence: manual	Medium
2.1.1	Keyboard	A	Partial	Some patterns; full functionality: manual	Critical
2.1.2	No Keyboard Trap	A	Partial	Common trap patterns; all paths: manual	Critical
2.1.4	Character Key Shortcuts	A	No	Cannot identify shortcuts	Low
2.2.1	Timing Adjustable	A	No	Cannot verify adjustment mechanisms	High
2.2.2	Pause, Stop, Hide	A	Partial	Blink/marquee; controls: manual	Medium

WCAG ID	Name	Level	Auto?	Detection Method	Priority
2.3.1	Three Flashes	A	Partial	Flash rate analysis; edge cases: manual	High
2.4.1	Bypass Blocks	A	Full	Skip link and landmark detection	High
2.4.2	Page Titled	A	Full	Title element presence (quality: manual)	Medium
2.4.3	Focus Order	A	Partial	Tabindex issues; logical sequence: manual	High
2.4.4	Link Purpose (Context)	A	Partial	Generic text flagging; context: manual	High
2.4.5	Multiple Ways	AA	No	Cannot verify navigation methods	Medium
2.4.6	Headings and Labels	AA	Partial	Empty headings; descriptiveness: manual	Medium
2.4.7	Focus Visible	AA	Partial	Focus indicator existence; all states: manual	High
2.5.1	Pointer Gestures	A	No	Cannot identify gestures	Medium
2.5.2	Pointer Cancellation	A	No	Cannot test down-event activation	Low
2.5.3	Label in Name	A	Full	Visible text vs. accessible name comparison	High
2.5.4	Motion Actuation	A	No	Cannot identify motion features	Low
3.1.1	Language of Page	A	Full	html lang attribute validation	Medium
3.1.2	Language of Parts	AA	Partial	lang attributes; unmarked changes: manual	Low
3.2.1	On Focus	A	No	Context change on focus: manual	Medium
3.2.2	On Input	A	Partial	Auto-submission detection	Medium
3.2.3	Consistent Navigation	AA	No	Cross-page comparison: manual	Medium
3.2.4	Consistent Identification	AA	No	Cross-page consistency: manual	Low
3.3.1	Error Identification	A	No	Error message presence/clarity: manual	High
3.3.2	Labels or Instructions	A	Partial	Missing labels; adequacy: manual	High
3.3.3	Error Suggestion	AA	No	Correction suggestions: manual	Medium
3.3.4	Error Prevention	AA	No	Confirmation mechanisms: manual	High
4.1.2	Name, Role, Value	A	Partial	Excellent ARIA coverage; semantics: manual	Critical
4.1.3	Status Messages	AA	No	ARIA live regions; announcements: manual	Medium

Summary: ~10 criteria fully automatable (~20%), ~22 partially automatable (~44%), ~18 require human review (~36%).

WCAG 2.2 Additions (October 2023)

WCAG 2.2 added 9 new success criteria and removed 4.1.1 Parsing (declared obsolete). The most impactful for automated detection: 2.5.8 Target Size Minimum (AA, fully automatable, measure element dimensions at least 24x24 CSS pixels) and 2.4.11 Focus Not Obscured (AA, partially automatable, detect overlap with sticky headers). New Level A criteria include 3.2.6 Consistent Help and 3.3.7 Redundant Entry, both requiring human review. The 3.3.8 Accessible Authentication criterion (AA) prohibits cognitive function tests like unassisted CAPTCHAs, a significant compliance consideration for login flows.

WCAG 3.0 Status and the APCA Question

WCAG 3.0 remains a Working Draft (latest: September 2025) and is not expected before 2028 at the earliest, with some estimates pushing to 2030. It introduces a graduated Bronze/Silver/Gold scoring model replacing pass/fail, and covers broader scope including apps, XR, voice, and wearables. The APCA (Advanced Perceptual Contrast Algorithm) will replace the current contrast ratio formula, fixing known problems where orange on white passes current WCAG but is hard to read.

> **Recommendation:** Build for WCAG 2.2 AA today. Optionally include APCA as a forward-looking feature. Design extensible architecture that can accommodate WCAG 3.0's scoring model when it stabilizes.

Section 508, EN 301 549, and the European Accessibility Act

Section 508 formally requires WCAG 2.0 Level AA for federal ICT and mandates VPAT documentation. ADA Title II now requires WCAG 2.1 AA. ADA Title III references no specific version but courts use 2.1 AA. Practical convergence: targeting WCAG 2.1 AA satisfies all three.

EN 301 549 (EU) incorporates WCAG 2.1 AA for web content and goes beyond WCAG for hardware, biometrics, and real-time communication. The European Accessibility Act took effect June 28, 2025, with penalties up to 4% of annual revenue in France and 100,000 euros per violation in Germany. It has extraterritorial reach: US companies serving EU consumers in e-commerce, banking, or digital services must comply.

3. Tool Comparison Matrix

Dimension	axe-core	WAVE API	Lighthouse	Pa11y	IBM Equal Access
Vendor	Deque Systems	WebAIM	Google	Community OSS	IBM
License	MPL-2.0 (free)	Free ext / paid API	Apache 2.0	LGPL-3.0	Apache 2.0
WCAG ver.	2.0, 2.1, 2.2	2.2 A/AA	2.0, 2.1 A/AA	2.0, 2.1 AA	2.2 A/AA
Rules	90+ rules	Proprietary set	~50+ audits	Varies by runner	Proprietary set
Detection	~57% real-world	~30-40%	Less than axe	Varies by engine	Comparable
False pos.	Zero FP design	More alerts	Binary pass/fail	Errors + Warnings	Violation/Review
API	JS/Node API, CLI	REST API (credits)	CLI + Node API	CLI + Node API	NPM package
Cost	Free (Pro $500+/yr)	$0.025/credit	Free	Free	Free
CI/CD	Excellent	Via API	Lighthouse CI	First-class	Via npm
Shadow DOM	Full (open)	No	Yes (via axe)	Limited	Most checks
Best for	Primary engine	Visual QA	Quick baseline	Site-wide CLI	Enterprise reports

Recommended stack: Primary engine: **axe-core via @axe-core/playwright**. This combination provides the industry-standard rule set (90+ rules, WCAG 2.2), the best browser automation framework (Playwright: cross-browser, auto-wait, Shadow DOM piercing, ARIA snapshots), and 2.5M+ weekly npm downloads confirming production maturity. Supplement with WAVE API for visual overlay data on critical pages and IBM Equal Access for cross-validation.

4. Agent Architecture Recommendation

Crawling Layer: Playwright (Chromium)

Playwright is the unambiguous choice over Puppeteer (Chrome-only, declining) and Selenium (legacy). Key advantages: native ARIA snapshot output (YAML accessibility tree), Shadow DOM piercing via Locator API, cross-browser testing (Chromium + Firefox + WebKit), built-in auto-wait that eliminates flakiness with SPAs, network interception for script injection, and browserContext.storageState for authenticated scanning.

Use a hybrid sitemap + link-following crawl strategy with configurable depth limits. Parse robots.txt by default but provide an --ignore-robots flag for first-party scanning. Rate-limit crawling to 1-2 seconds between requests with adaptive adjustment. Limit to 2-3 concurrent connections per domain.

Detection Layer: axe-core + CDP Accessibility Tree

Run @axe-core/playwright with tags [wcag2a, wcag2aa, wcag21a, wcag21aa, wcag22aa] on each rendered page. Simultaneously extract the full accessibility tree via Chrome DevTools Protocol (Accessibility.getFullAXTree). The axe-core results provide structured violation data with impact levels; the CDP tree provides the semantic structure for AI analysis.

AI Analysis Layer: LLM Integration

This is where an AI agent differentiates from existing tools. Pass axe-core's "needs review" items plus the accessibility tree context to an LLM for:

- Alt text quality evaluation (not just presence)
- Heading hierarchy semantic assessment
- ARIA pattern correctness beyond syntax
- Color context interpretation
- Natural-language remediation guidance
- VPAT pre-population from scan data

This could push practical coverage from 57% toward 70-80%, a meaningful improvement over any existing tool.

SPA-Specific Handling

SPAs (React, Vue, Angular) require special treatment because client-side routing doesn't trigger full page loads, so screen readers receive no navigation signal. The scanner must simulate navigation events and wait for route transitions, test focus management after navigation, detect ARIA live regions, and check that document.title updates per route. Pages using ARIA had 57 errors on average versus 27 for pages without ARIA (WebAIM 2025), confirming that misused ARIA is worse than none.

Report Output Formats

The agent should generate five output types: JSON (developer/API consumption), HTML (interactive drill-down with screenshots), PDF (executive/legal/compliance), CSV (issue tracking import), and VPAT/ACR (enterprise procurement). VPATs use the ITI template with four editions (508, EU, WCAG, International). An AI layer can pre-populate VPAT entries from scan data, reducing the $1,000-5,000 manual cost to a review-and-confirm

workflow.

5. Violation Taxonomy

The "Big Six": 96% of Automatically Detected Errors

These six violation categories account for the vast majority of detected issues and should be the agent's highest priority.

Violation	WCAG	Sites Affected	Avg/Page	Auto?	Legal Risk
Missing/poor alt text	1.1.1 (A)	55.5%	11	Partial	Very High
Insufficient contrast	1.4.3 (AA)	79.1%	29.6	Full	Very High
Empty links/buttons	2.4.4 / 4.1.2	45.4% / 29.6%	N/A	Full	High
Missing form labels	1.3.1 / 4.1.2	48.2%	N/A	Full	Very High
Missing doc language	3.1.1 (A)	15.8%	1	Full	Medium
Skipped headings	1.3.1 (A)	39%	N/A	Full	Low-Med

Missing/Poor Alt Text (1.1.1, Level A)

Automated detection identifies missing alt attributes and empty alt on linked images. AI can evaluate whether existing alt text is meaningful, flagging generic patterns like alt="image" or filename-based text. Of missing alt text instances, 44% involve linked images, making them navigation-breaking. Fix: add descriptive alt for informative images, alt="" for decorative images, and alt describing the link destination for linked images.

Insufficient Color Contrast (1.4.3 / 1.4.11, Level AA)

The single most common violation, affecting 79.1% of homepages with an average of 29.6 instances per page. Required ratios: 4.5:1 for normal text (under 18pt), 3:1 for large text (18pt+ or 14pt+ bold), 3:1 for non-text UI components. Fully automatable via luminance ratio calculation, though text-over-image/gradient scenarios require rendered-page analysis. The minimum passing gray on white is #767676.

Empty Links and Buttons (2.4.4 / 4.1.2, Level A)

Appear on 45.4% and 29.6% of homepages respectively. Fully automatable via accessible name computation. Common causes: icon buttons without aria-label, image links without alt text, links wrapping only whitespace. Fix: add aria-label to icon buttons, alt text to linked images, visible or screen-reader text to empty elements.

Missing Form Input Labels (1.3.1 / 4.1.2, Level A)

Affect 48.2% of homepages. Automatable via programmatic label association verification (for/id pairing, wrapping label, or aria-label/aria-labelledby). Fix priority: explicit label for="" (preferred), aria-label (when design requires hidden label), aria-labelledby (when label text exists elsewhere). Group related fields with fieldset and legend.

Additional High-Frequency Violations

Keyboard Navigation Failures (2.1.1 / 2.1.2 / 2.4.7)

Primarily manual detection. Automated tools can detect tabindex misuse, div/span elements with click handlers but no keyboard support, and CSS outline:none without replacement focus styles. The AI layer should simulate tab navigation and flag elements unreachable by keyboard. Fix: use semantic HTML (button not div onclick), add onkeydown handlers for custom widgets, implement focus trap for modals with Escape key dismissal. Legal risk: Very high.

Videos Without Captions (1.2.2, Level A)

Partially automatable: detect video elements without track kind="captions". Caption quality requires manual review. Professional captioning runs $1-3/minute of video. Legal risk: High.

Inaccessible PDFs

WCAG applies to all web content including PDFs. Only 3.2% of academic PDFs meet full accessibility requirements. Automated detection can identify untagged PDFs and missing structure. Manual remediation costs ~$3-4/page. Tools: Adobe Acrobat Pro (manual), CommonLook (enterprise), Equidox (AI-powered). Legal risk: High.

6. Legal Risk Map

Violations Most Likely to Trigger Litigation

Based on lawsuit filing data from UsableNet, Seyfarth Shaw, and settlement reports, ranked by frequency:

Rank	Violation	WCAG	Legal Risk	Notes
1	Missing alt text on images	1.1.1	Critical	Cited in virtually every ADA complaint
2	Screen reader incompatibility	4.1.2 + ARIA	Critical	Missing accessible names, broken roles
3	Keyboard nav failures	2.1.1, 2.1.2	Critical	Transaction completion barriers
4	Missing form labels	1.3.1, 4.1.2	Very High	48.2% of sites affected
5	Insufficient contrast	1.4.3	Very High	79.1% of sites; sheer volume
6	Missing video captions	1.2.2	High	Growing as video becomes ubiquitous
7	Empty links/buttons	2.4.4, 4.1.2	High	Navigation barriers

Industry Targeting Patterns

77%
Lawsuits target e-commerce

77%
Target sub-$25M revenue

746%
Illinois filing growth (2025)

E-commerce and retail account for 77% of ADA website lawsuits. Serial plaintiffs deliberately target small businesses with standardized demand letters. Industries most at risk: e-commerce, food and beverage, restaurants, apparel, healthcare, hospitality, and education. New York, California, and Florida account for the majority of filings, with Illinois seeing 746% growth in 2025 as plaintiff firms migrated from NY federal courts that tightened standing requirements.

Settlement and Defense Costs

Typical settlements: $5,000 to $75,000 plus attorney fees and mandatory website remediation. For every filed lawsuit, approximately 10 demand letters are sent, typically settling for $5,000 to $20,000. Legal defense costs add $10,000 to $50,000+. Total cost including remediation ranges from $20,000 to $200,000+ depending on site complexity.

Key finding: 25% of 2024 lawsuits targeted sites already using overlay widgets. Courts have interpreted overlay installation as documenting awareness of accessibility issues while choosing a cosmetic fix, potentially establishing willful negligence.

The Circuit Split

The Supreme Court has repeatedly declined to resolve whether websites are "places of public accommodation" under ADA Title III. The broad view (1st, 2nd, 7th Circuits) holds websites can independently be public

accommodations. The nexus theory (3rd, 6th, 9th Circuits) requires connection to a physical location. The 11th Circuit's Gil v. Winn-Dixie opinion (holding websites are not public accommodations) was vacated in December 2021, leaving no binding precedent. A 2025 Minnesota district court ruled that web-only businesses are subject to Title III, potentially expanding scope. Until the Supreme Court takes up this question, the legal landscape remains jurisdiction-dependent.

7. Open Questions and Human Judgment Requirements

What the Agent Cannot Auto-Detect (the 36% Gap)

The following WCAG criteria fundamentally cannot be automated and must be flagged as "requires human review" in any agent output:

- **Meaningful sequence** (1.3.2): whether reading order preserves meaning when CSS is disabled
- **Sensory characteristics** (1.3.3): instructions relying solely on shape, color, size, location, or sound
- **Content on hover/focus** (1.4.13): whether hoverable content is dismissible, persistent, and hoverable
- **Timing adjustable** (2.2.1): whether time limits can be turned off, adjusted, or extended
- **Multiple ways** (2.4.5): whether multiple navigation methods exist
- **Consistent navigation** (3.2.3): whether nav appears in same relative order across pages
- **Error identification/suggestion** (3.3.1, 3.3.3): whether error messages exist and are helpful
- **Error prevention** (3.3.4): whether submissions are reversible, checked, or confirmed
- **Status messages** (4.1.3): whether ARIA live regions correctly announce dynamic content
- **All cognitive accessibility criteria**: readability, clarity, and cognitive load assessment

Where AI Can Push the Boundary

An LLM-augmented agent can meaningfully address the "partially automatable" middle tier (~44% of criteria):

- **Alt text quality assessment:** evaluate whether existing alt text meaningfully describes image content in context
- **Heading hierarchy semantics:** beyond checking for skipped levels, assess whether structure logically represents content
- **ARIA pattern correctness:** determine whether custom ARIA implementations follow WAI-ARIA Authoring Practices semantically
- **Link purpose evaluation:** NLP assessment of whether link text provides sufficient context
- **Focus order logic:** combining tab order extraction with page layout analysis to flag illogical sequences
- **VPAT pre-population:** generating draft Accessibility Conformance Reports from scan data

Market Positioning

The digital accessibility software market reached $0.80 billion in 2025, projected to $1.08 billion by 2030 (6.31% CAGR). The three biggest gaps an AI-native agent could fill:

- **Source-code remediation at SMB price points** ($5K-15K/year versus $100-500/page manual audits)
- **Context-aware alt text and semantic understanding** beyond binary presence checking
- **Continuous compliance with natural-language developer guidance** embedded in CI/CD workflows

Enterprise buyers expect VPATs, continuous monitoring, and Jira integration. SMBs need affordable, real fixes. The 77% of lawsuits targeting sub-$25M-revenue companies represents an underserved market trapped between useless overlays and unaffordable enterprise solutions.

RECENT REGULATORY CHANGES (LAST 18 MONTHS)

- DOJ rescinded 11 ADA guidance documents (March 2025) and will not pursue pending rulemakings
- DOJ announced review of all ADA Title II and III regulations (October 2025), no timetable
- Title II compliance deadline of April 24, 2026 is imminent
- FTC fined AccessiBe $1M (April 2025): first federal regulatory action against an overlay vendor
- European Accessibility Act enforcement began June 28, 2025 with extraterritorial reach to US companies
- ADA lawsuit volume surged: 5,114 total in 2025, H1 2025 up 37% YoY, Illinois filings up 746%
- Pro se AI-assisted filings increased 40% in 2025
- WCAG 2.2 published October 2023: 9 new criteria, 4.1.1 Parsing declared obsolete
- WCAG 3.0 Working Draft updated September 2025, finalization not expected before 2028
- WebAIM Million 2026: accessibility regression, 95.9% failure rate (up from 94.8%), 56.1 errors/page