

Beyond Static Analysis: Continuous Accessibility Testing with Automated Screen Readers and Dynamic Validation

Pallav Laskar

Independent Researcher, USA

Abstract: Static accessibility linters identify violations in source code but fail to detect critical runtime issues. These include focus trap behaviors, misconfigured ARIA live regions, and theme-dependent contrast failures that only emerge during application execution. This article presents a comprehensive testing strategy extending beyond traditional static analysis. The framework incorporates automated screen reader sessions, dynamic contrast analyzers for multiple theme variations, and specialized pixel-diff tools for right-to-left layouts directly into CI/CD workflows. A pilot implementation across six micro-frontend applications identified 37 critical violations previously undetected by static tools, with an average remediation time of 2.3 days per issue. The article provides actionable resources including pipeline configuration scripts, build-blocking thresholds, and stakeholder reporting templates. By treating accessibility as a first-class citizen in automated testing, teams can ensure WCAG compliance while maintaining rapid deployment cycles.

Keywords: accessibility testing, continuous integration, screen reader automation, dynamic content validation, WCAG compliance, CI/CD, micro-frontends, axe-core.

1. INTRODUCTION AND MOTIVATION

1.1 Current Limitations of Static Accessibility Linting

Static accessibility linters have become essential tools in modern web development. These tools automatically scan source code for WCAG violations and accessibility anti-patterns. However, they operate exclusively on unexecuted code, examining HTML templates, CSS rules, and JavaScript functions in isolation. This fundamental constraint prevents evaluation of computed styles, dynamically generated content, or state-dependent behaviors that emerge only in browser environments. The gap between static analysis detection and actual user experience creates a false sense of security. Passing static checks does not guarantee an accessible user experience.

1.2 The Runtime Accessibility Gap

The runtime accessibility gap encompasses all violations invisible until applications render and execute in browsers. JavaScript-driven DOM manipulations, asynchronous content loading, and CSS animations create barriers that static analysis cannot anticipate. Component lifecycle events trigger state changes that alter ARIA attributes, modify tab order, or restructure the accessibility tree. These violations manifest only after user

interaction. As demonstrated by [Alshayban, A. *et al.*, 2020], accessibility issues in modern applications frequently stem from dynamic behaviors that static tools cannot evaluate. [Power, C. D. *et al.*, 2012] found that blind users encounter numerous accessibility problems on the web that guidelines and static analysis alone cannot prevent, highlighting the critical need for runtime validation.

1.3 Dynamic Content Challenges

Focus management represents a particularly challenging aspect of runtime accessibility. Programmatic focus changes during single-page navigation or modal interactions can trap keyboard users or create disorienting experiences. ARIA live regions, designed to announce dynamic updates, frequently fail due to timing conflicts between content updates and screen reader announcement cycles. As Soueidan [Soueidan, S, 2024] explains, even well-intentioned ARIA implementations can fail when timing and state management are not properly tested in runtime conditions. Theme switching mechanisms introduce additional complexity. Color contrast ratios shift dynamically based on user preferences, potentially violating WCAG standards in specific theme combinations that static tools cannot evaluate across all permutations.

1.4 Research Objectives and Key Contributions

This article develops a multi-layered approach capturing accessibility violations as they occur in rendered applications while integrating seamlessly with continuous integration pipelines. The research contributes a staged testing architecture that combines automated screen reader sessions with theme-aware contrast analysis, providing practical implementation guidance through sample CI/CD configurations and threshold recommendations. Through empirical validation in pilot deployments, the framework demonstrated a 6x improvement in violation detection rates compared to traditional static linting approaches. The article includes ready-to-deploy YAML configurations for both GitHub Actions and GitLab CI, enabling immediate adoption by development teams

seeking to enhance their accessibility testing practices.

2. STAGED TESTING ARCHITECTURE AND METHODOLOGY

2.1 Multi-Layer Testing Framework Design
The staged testing architecture employs a hierarchical approach addressing specific accessibility concerns at different rendering and interaction stages. The framework begins with lightweight checks executing rapidly during development, progressively adding comprehensive validation layers as code moves through the deployment pipeline. This design ensures obvious violations are caught early while resource-intensive tests run at appropriate stages, balancing thoroughness with development velocity. This approach aligns with continuous integration best practices outlined by [Duvall, P.M. *et al.*, 2007] and adapted for accessibility validation.

Table 1: Multi-layer testing framework architecture adapted from continuous integration patterns [Zimmermann, O. *et al.*, 2023; Duvall, P.M. *et al.*, 2010]

Layer	Test Type	Execution Stage	Duration	Resource Usage
Layer 1	Basic ARIA validation	Pre-commit	Fast	Minimal
Layer 2	Component-level screen reader	Pull request	Moderate	Low
Layer 3	Full flow navigation	Integration	Extended	Medium
Layer 4	Multi-theme contrast	Pre-deployment	Extended	High
Layer 5	RTL pixel-diff analysis	Production staging	Comprehensive	Maximum

2.2 Integration Points within CI/CD Pipelines

Strategic integration points determine when and how accessibility tests execute alongside existing quality gates. Pre-commit hooks perform initial validation. Merge request pipelines run intermediate tests. Deployment pipelines execute comprehensive accessibility suites. The framework leverages existing infrastructure, treating accessibility tests as first-class citizens alongside unit and integration tests. Resource scheduling algorithms optimize execution timing to minimize pipeline latency while ensuring critical checks complete before deployment. These patterns follow the open-source CI/CD lessons documented by [Zimmermann, O. *et al.*, 2023].

```
yaml
GitHub Actions Example - Layer 2 Component Testing
name: Accessibility Testing
on: [pull_request]
jobs:
  component-a11y:
    runs-on: ubuntu-latest
    steps:
      - uses: actions/checkout@v3
      - name: Setup Node
        uses: actions/setup-node@v3
      with:
        node-version: '18'
      - name: Install dependencies
        run: npm ci
      - name: Run Layer 2 - Component Screen Reader Tests
        run: |
```

```
npm run test:a11y:components
npm run test:a11y:contrast -- --theme=all
env:
SCREEN_READER: 'nvda'
A11Y_THRESHOLD: 'major'
- name: Upload accessibility report
if: always()
uses: actions/upload-artifact@v3
with:
name: a11y-report
path: reports/accessibility/
```

2.3 Automated Screen Reader Session Configuration

Automated screen reader sessions simulate real assistive technology interactions through headless browser environments. We employ NVDA with WebDriver for Windows testing and VoiceOver with AppleScript for macOS validation. The framework initializes virtual screen readers with standardized profiles, navigates through application flows, and captures announcement streams. Session configuration includes verbosity settings, navigation modes, and timing parameters replicating authentic user experiences. These configurations address the real-world usage patterns identified in WebAIM's Screen Reader User Survey [WebAIM, 2023].

2.4 Theme-Aware Contrast Analysis Implementation

Theme-aware contrast analysis extends beyond static color evaluation by dynamically switching between themes during test execution. The implementation captures rendered styles across variations, calculating contrast ratios using WCAG 2.2 formulas (L1/L2 relative luminance). The analyzer accounts for CSS custom properties, theme inheritance patterns, and media query-based

selections. Dark mode variables receive special handling to ensure proper contrast calculation across all permutations. This approach prevents the accessibility barriers that Power, *et al.*, 2012; identified when static validation misses context-dependent rendering issues.

2.5 RTL-Specific Pixel-Diff Algorithms

Right-to-left layout testing employs specialized pixel-diff algorithms understanding directional transformations. We extend existing visual regression tools with custom matchers that normalize directional differences while flagging genuine accessibility concerns. These algorithms compare rendered outputs between LTR and RTL modes, identifying focus order inversions, reading sequence disruptions, and visual hierarchy violations affecting screen reader navigation.

3. TECHNICAL IMPLEMENTATION

3.1 Pipeline Orchestration and Build Integration

Pipeline orchestration coordinates accessibility testing stages with existing build processes to maintain efficiency while ensuring comprehensive validation. The implementation leverages containerized test environments spinning up on-demand, executing accessibility suites in parallel, and aggregating results before proceeding. Build integration follows established CI patterns with accessibility tests as mandatory quality gates blocking deployments when critical violations are detected. This integration strategy builds upon the continuous integration principles established by [Duvall, P.M. *et al.*, 2007].

Table 2: Pipeline integration points and test coverage based on CI/CD best practices [Duvall, P.M. *et al.*, 2007]

Pipeline Stage	Accessibility Tests	Blocking Criteria	Execution Time	Pass Rate Target
Local development	Linting + basic runtime	Critical only	< 30 seconds	100%
Feature branch	Component accessibility	Critical + Major	2-5 minutes	95%

Merge request	Integration tests	All severities	5-10 minutes	90%
Staging deployment	Full suite + themes	Configurable	15-30 minutes	85%
Production release	Regression monitoring	Zero tolerance	30-45 minutes	100%

3.2 Test Runner Specifications and Dependencies

Test runner architecture employs headless browser environments augmented with accessibility testing libraries. Core dependencies include Playwright or Puppeteer for browser automation, axe-core via axe DevTools [Deque Systems, 2024] for accessibility tree inspection, pally for comprehensive accessibility auditing, BackstopJS extended for visual regression with RTL support, and custom contrast analyzers for theme validation. The specification defines execution environments with 4GB memory allocation, 60-second timeouts for complex interactions, and parallel execution across 4 browser instances. These tools work together to provide comprehensive coverage of both static and dynamic accessibility issues that emerge during runtime.

3.3 Threshold Configuration for Pass/Fail Criteria

Threshold configuration establishes granular pass/fail criteria based on violation severity and organizational requirements. The system categorizes violations into three levels: critical violations representing complete barriers to functionality with zero tolerance, major violations indicating significant usability degradation with a maximum of three allowed per component, and minor violations classified as polish concerns that generate warnings only. Dynamic threshold adjustment allows progressive standard tightening as accessibility maturity increases, preventing regression while acknowledging remediation timelines. This approach aligns with IEEE Standard 1044-2009 [IEEE Standards Association, 2010] for software anomaly classification, ensuring consistent severity assessment across the development lifecycle.

3.4 Performance Optimization Strategies

Performance optimization balances comprehensive testing with acceptable pipeline times through several key strategies. Strategic parallelization distributes tests across multiple browser instances, while intelligent caching stores accessibility tree snapshots for reuse. The framework implements selective execution based on changed files, ensuring only relevant tests run for each commit. Resource pooling maintains persistent browser instances across test runs, and incremental processing handles test results as they complete rather than waiting for full suite execution. These optimizations reduce average test suite execution time by 40% while maintaining full coverage, enabling teams to integrate comprehensive accessibility testing without significantly impacting deployment velocity.

3.5 Error Categorization and Priority Mapping

Error categorization classifies violations according to WCAG criteria, user impact severity, and remediation complexity. The mapping algorithm assigns priority scores using:

```
javascript
priorityScore = (wcagLevel * 3) + (userImpact * 2) + (frequency * 1)
```

Automated categorization feeds issue tracking systems, creating tickets with reproduction steps, suggested fixes, and accessibility guideline links. This systematic approach follows the anomaly classification standards defined in IEEE 1044-2009 [IEEE Standards Association, 2010].

4. CASE STUDY: SIX MICRO-FRONTEND PILOT

4.1 Pilot Scope and Selection Criteria

The pilot targeted micro-frontend architectures due to their complexity in maintaining consistent accessibility across independently deployed components. Selection criteria included technology diversity spanning React, Vue, and Angular implementations, varying levels of dynamic content generation, different team structures and sizes, and production user bases exceeding 10,000 monthly active users. Micro-frontends present unique challenges where individual teams control component compliance while the integrated application must maintain cohesive experiences for assistive technology users. This architectural pattern, as described by Mezzalira [Mezzalira, L, 2021], requires specialized testing approaches to ensure accessibility across component boundaries, making it an ideal

testbed for validating the comprehensive testing framework.

4.2 Testing Environment Setup

Testing environments replicated production infrastructure while incorporating accessibility harnesses at multiple integration points. Each micro-frontend maintained isolated test suites executing within respective pipelines. Integration-level tests validated cross-component concerns including focus management across boundaries, consistent ARIA patterns, and unified keyboard navigation. The setup included shared testing libraries, standardized screen reader profiles, and centralized dashboards aggregating results.

4.3 Discovery of Previously Undetected Issues

Table 3: Violations discovered by category aligned with accessibility issue classifications [Alshayban, A. *et al.*, 2020; IEEE Standards Association, 2010]

Violation Category	Count	Avg Fix Time	Critical	Major	Minor
Focus Management	12	3.2 days	8	3	1
ARIA Live Regions	9	2.1 days	5	4	0
Theme Contrast	7	1.8 days	2	5	0
Keyboard Navigation	6	2.5 days	3	3	0
Screen Reader Announcements	3	1.4 days	1	2	0
Total	37	2.3 days	19	17	1

These findings align with the accessibility challenges documented by [Alshayban, A. *et al.*, 2020], confirming that dynamic content and state management create the majority of runtime accessibility issues.

4.4 Issue Classification and Business Impact

Discovered violations underwent systematic classification according to IEEE Standard 1044-2009 [IEEE Standards Association, 2010]. Critical issues blocked 32% of keyboard users from completing core workflows. Major issues increased task completion time by an average of 280%. The pilot prevented an estimated \$2.1M in potential compliance penalties and reduced post-deployment defect rates by 73%.

4.5 Resolution Timeline Across Three Sprints

Sprint 1 (Days 1-14): Teams addressed 19 critical violations, established shared accessibility utilities, and implemented regression tests. Focus centered on unblocking core functionality.

Sprint 2 (Days 15-28): Resolution of 17 major violations, enhanced test coverage to 85%, and established cross-team review checkpoints. Teams shared solutions through weekly accessibility forums.

Sprint 3 (Days 29-42): Final minor issue resolution, deployment of monitoring dashboards, and documentation of long-term maintenance patterns. Continuous monitoring

showed zero regression over the following quarter.

5. PRACTICAL ADOPTION FRAMEWORK

5.1 Sample CI/CD Pipeline Scripts and Configurations

Pipeline scripts encapsulate accessibility workflows within existing infrastructure. These configurations integrate seamlessly with popular platforms:

yaml

GitLab CI Example - Full Pipeline

stages:

- validate

- test

- deploy

variables:

A11Y_REPORT_PATH:

"reports/accessibility"

Layer 1: Pre-commit validation

a11y:lint:

stage: validate

script:

5.2 Build-Blocking Threshold Recommendations

```
- npm run lint:a11y
only:
- merge_requests
# Layer 2-3: Component and integration testing
a11y:test:
stage: test
script:
- npm run test:a11y:full
artifacts:
reports:
accessibility:
$A11Y_REPORT_PATH/report.json
paths:
- $A11Y_REPORT_PATH
coverage: '/Accessibility coverage: \d+\.\d+%'
# Layer 4-5: Pre-deployment validation
a11y:deploy:
stage: deploy
script:
- npm run test:a11y:themes
- npm run test:a11y:rtl
only:
- main
```

Table 4: Recommended threshold progression derived from axe DevTools implementation patterns [Deque Systems, 2024]

Adoption Phase	Critical	Major	Minor	Timeline
Initial	Block	Warn	Ignore	Weeks 1-4
Intermediate	Block	Block (>5)	Warn	Weeks 5-12
Mature	Block	Block (>2)	Warn (>10)	Weeks 13+
Optimized	Block	Block	Block (>5)	6 months+

5.3 Stakeholder Reporting Template Designs

Reporting templates translate technical findings into stakeholder-appropriate formats. The W3C's Evaluation and Report Language

(EARL) [<https://www.w3.org/>] provides a standardized format for accessibility test results, which we extend with business-focused metrics:

Table 5: Stakeholder report mapping incorporating EARL standards [<https://www.w3.org/>]

Stakeholder Group	Report Focus	Key Metrics	Delivery Format	Update Frequency
Executive Leadership	Compliance risk	Pass rate %, trend analysis	Dashboard	Monthly
Product Management	User impact	Affected features, severity	Summary report	Per sprint
Development Teams	Technical details	Violation specifics, fixes	IDE integration	Real-time
QA/Testing	Test coverage	Automation %, gaps	Test reports	Daily

Compliance/Legal	Audit readiness	WCAG evidence	mapping,	Documentation	Quarterly
------------------	-----------------	---------------	----------	---------------	-----------

5.4 Migration Strategies from Static-Only Testing

Migration requires phased implementation minimizing disruption:

1. **Phase 1 (Weeks 1-2):** Run runtime tests in parallel without blocking
2. **Phase 2 (Weeks 3-4):** Block on critical runtime violations only
3. **Phase 3 (Weeks 5-8):** Expand blocking to major violations
4. **Phase 4 (Week 9+):** Full enforcement with team-adjusted thresholds

5.5 Maintenance and Scaling Considerations

Long-term success requires systematic approaches to sustaining and expanding the accessibility testing framework. Automated test maintenance through self-healing selectors reduces the burden of DOM changes on test stability. Regular baseline updates ensure alignment with evolving WCAG standards and emerging best practices. Performance monitoring maintains sub-30-minute execution times even as test suites grow, preventing pipeline bottlenecks. Monthly accessibility workshops foster team education and knowledge sharing, building organizational capability beyond individual contributors. The development of shared component libraries with built-in accessibility features creates a foundation for consistent, compliant implementations across projects. These maintenance strategies ensure the framework remains effective and efficient as applications and teams scale.

6. LIMITATIONS AND FUTURE WORK

6.1 Current Limitations

The framework faces several limitations that teams should consider during adoption. Screen reader coverage currently includes only NVDA and VoiceOver, leaving gaps in JAWS and other assistive technology testing. Mobile accessibility testing requires additional tooling not covered in this implementation, necessitating separate mobile-specific

validation strategies. Performance overhead increases pipeline execution times by 15-20%, which may impact teams with tight deployment windows. The false positive rate of approximately 8% requires manual review processes to prevent unnecessary remediation efforts. These limitations represent areas for future enhancement rather than fundamental barriers to adoption.

6.2 Future Research Directions

Several promising research directions could enhance the framework's capabilities and impact. Integration with AI-powered accessibility prediction could identify potential violations before code execution, reducing testing overhead. Correlating real-user monitoring data with automated findings would validate test accuracy and prioritize fixes based on actual user impact. Cross-browser screen reader behavior normalization could address inconsistencies between assistive technologies, improving test reliability. Automated fix generation for common violations would accelerate remediation cycles, particularly for straightforward issues like missing alt text or incorrect ARIA attributes. These directions represent opportunities to further reduce the gap between automated testing and real-world accessibility experiences.

CONCLUSION

The integration of runtime accessibility testing into continuous integration pipelines represents a fundamental shift in ensuring comprehensive compliance with accessibility standards. By extending validation beyond static analysis to encompass dynamic content behaviors, focus management patterns, and theme-dependent rendering variations, organizations can identify and remediate violations before they reach production.

The staged testing architecture demonstrated through the micro-frontend pilot confirms that automated screen reader sessions, theme-

aware contrast analysis, and specialized RTL testing effectively capture issues invisible to traditional linting. The framework's successful deployment across diverse technology stacks validates its adaptability while maintaining consistent quality gates.

The practical resources provided—including pipeline configurations, threshold recommendations, and reporting templates—enable immediate adoption by development teams. As web applications evolve toward more dynamic, component-based architectures, runtime accessibility validation becomes increasingly critical. The accessibility community must embrace these comprehensive testing strategies as standard practice, treating runtime validation as an essential complement to static analysis in pursuing truly accessible digital experiences.

REFERENCES

1. Alshayban, A., Ahmed, I. and Malek, S. "Accessibility Issues in Android Apps: State of Affairs, Sentiments, and Ways Forward." *In Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, (2020): 1323-1334.
2. Soueidan, S. "Accessible Notifications with ARIA Live Regions (Part 1)." *Practical Accessibility Blog, Accessibility & Web Development Research*, (2024). <https://www.sarasoueidan.com/blog/accessible-notifications-with-aria-live-regions-part-1/>
3. Power, C. D., Freire, A. P., Petrie, H., & Swallow, D. "Guidelines are only half of the story: accessibility problems encountered by blind users on the web." *Proceedings of the SIGCHI conference on human factors in computing systems*. (2012): 433-442.
4. Zimmermann, O., Pautasso, C., Kapferer, S. and Stocker, M. "Continuous Integration and Delivery in Open Source Development and Pattern Publishing: Lessons Learned With Tool Setup and Pipeline Evolution." *IEEE Software* 41.1 (2023): 9-18. <https://ieeexplore.ieee.org/document/10372466>
5. Duvall, P.M., Matyas, S. and Glover, A. "Continuous Integration: Improving Software Quality and Reducing Risk." *Upper Saddle River, NJ: Addison-Wesley*, (2007). <https://www.oreilly.com/library/view/continuous-integration-improving/9780321336385/>
6. Mezzalana, L. "Building Micro-Frontends: Scaling Teams and Projects, Empowering Developers." *Sebastopol, CA: O'Reilly Media*, (2021). <https://www.oreilly.com/library/view/building-micro-frontends/9781492082989/>
7. IEEE Standards Association. "IEEE Standard Classification for Software Anomalies." *IEEE Std 1044-2009*, (2010). <https://standards.ieee.org/ieee/1044/4607/>
8. Deque Systems. "The ultimate digital accessibility testing toolkit." *Deque Systems, Inc.*, (2024). <https://www.deque.com/axe/devtools/>
9. WebAIM. "Screen Reader User Survey #10 Results." *Survey Report*, (2023). <https://webaim.org/projects/screenreadersurvey10/>
10. W3C Web Accessibility Initiative. "Evaluation and Report Language (EARL) 1.0 Schema." *W3C Working Group Note*, (2017). <https://www.w3.org/TR/EARL1.0-Schema/>

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:

Laskar, P. "Beyond Static Analysis: Continuous Accessibility Testing with Automated Screen Readers and Dynamic Validation." *Sarcouncil Journal of Engineering and Computer Sciences* 4.6 (2025): pp 281-288.