Automated ACR Evaluation versus Manual Desk Audit

A Feature-by-Feature Comparison of Phase 1 ACR Review Methods for ICT Procurement

An evidence-based comparison of automated, rubric-based ACR evaluation and expert-driven manual desk audit for document-level (Phase 1) accessibility review.

DOCUMENT CONTROL

SERIES
ACR Evaluator Technical Series

REPORT TYPE
Comparative analysis

VERSION
v1 · June 2026

CLASSIFICATION
Public

SCOPE
Phase 1 — document-level ACR review

ABSTRACT

Public-sector buyers frequently rely on vendor Accessibility Conformance Reports (ACRs), often prepared using the Voluntary Product Accessibility Template (VPAT), to inform ICT procurement, acceptance, renewal, and oversight decisions. Traditionally, these reports are reviewed through an expert-driven manual desk audit, in which an accessibility professional evaluates the completeness, credibility, and internal consistency of the vendor's claims.

This report compares that manual model with Revelo Software's ACR Evaluator, which uses agentic AI, robotic process automation, expert-prepared rubrics, rules-based applicability logic, structured scoring, and human-in-the-loop escalation. The analysis examines cost, turnaround time, repeatability, scalability, evidence confidence, limitations, and appropriate use cases.

INDEPENDENCE & METHODOLOGY STATEMENT

The findings in this report are derived from a defined evidence base and from explicit, stated assumptions rather than from marketing claims. The methodology — the comparison dimensions, the cost and turnaround inputs, and the evaluation boundary — is disclosed in full so that any reader can reproduce or challenge the analysis. **Disclosure:** the ACR Evaluator is a product of Revelo Software, and the commercial price and turnaround figures used here were specified for this comparison. The analysis is limited to Phase 1, document-level ACR review and does not assert that either method, on its own, establishes actual product conformance.

EXECUTIVE SUMMARY

Manual ACR desk audits remain useful because they apply expert judgment to vendor documentation. Experienced reviewers can identify incomplete ACRs, vague testing methods, unsupported conformance claims, outdated reports, inconsistent remarks, and misuse of *Not Applicable*. However, desk audits are labor-intensive, relatively expensive, difficult to scale, and inherently limited by the quality of the vendor-provided document. They generally do not verify actual product accessibility unless paired with hands-on testing.

The ACR Evaluator changes the economics and repeatability of Phase 1 ACR review. At a commercial price of $39 per report, compared with approximately $1,500 for an expert-driven manual desk audit and report, the ACR Evaluator provides a lower-cost first-pass review model. It also delivers results in up to 10 minutes, compared with up to two weeks for manual desk audit delivery. This speed and price differential makes it practical to screen more vendor ACRs earlier in the procurement lifecycle.[1]

The ACR Evaluator is not designed to remove expert judgment from high-risk decisions. Instead, it uses human-in-the-loop escalation. When the Evidence Confidence Score (ECS) is sufficiently low, when documentation is incomplete or contradictory, or when other risk indicators are elevated, the system triggers expert intervention. This creates a tiered model: automation handles structured, repeatable review tasks, while human experts focus on cases where judgment, validation, and remediation guidance are most needed.

> **CORE FINDING**
> The most effective operating model is not automation instead of experts. It is automation plus experts: use the ACR Evaluator for standardized, low-cost, rapid Phase 1 screening; use ECS and risk indicators to route low-confidence or high-impact cases to expert review and, when warranted, independent verification and validation.

EVIDENCE BASE AND SCOPE

This report compares two approaches to Phase 1 review of vendor ACRs: expert-driven manual document review and the ACR Evaluator automated review workflow. The comparison uses the following evidence base:

- The ACR Evaluator Technical Series white paper describing the rules-based applicability engine, ICT Classification-to-Criteria Crosswalk Framework, and Accessibility Evaluation Framework.
- The Manual Document Review / Desk Audit working assumptions, including cost, turnaround time, and common weaknesses of vendor ACR desk audits.
- User-specified commercial assumptions: $39 per ACR Evaluator report, approximately $1,500 per expert-driven manual desk audit and report, up to 10 minutes for ACR Evaluator delivery, and up to two weeks for manual desk audit delivery.

[1] Pricing note: the $39 ACR Evaluator report price is the commercial price assumption used for this comparison. The ACR Evaluator is offered with enterprise and volume discounts to public-sector agencies; actual pricing may vary by contract vehicle, purchase volume, licensing model, reseller terms, and agency requirements.

BOUNDARY OF ANALYSIS

This report evaluates Phase 1 document-level ACR evaluation. It does not claim that either method, by itself, proves actual product conformance. Hands-on testing, assistive technology validation, and user testing remain necessary when evidence is weak, risk is high, or procurement consequences warrant independent validation.

1. BACKGROUND: THE ROLE OF ACR REVIEW IN ICT PROCUREMENT

An Accessibility Conformance Report is intended to help buyers understand how an ICT product or service conforms to accessibility standards such as Section 508, WCAG, Section 255, and EN 301 549. In practice, ACRs vary widely in quality. Some include detailed testing methods, current version information, clear remarks, and meaningful limitations. Others contain vague claims, outdated product references, incomplete criteria, unsupported conformance levels, or inconsistent explanations.

For public-sector procurement teams, the challenge is not only whether a vendor has submitted an ACR. The larger question is whether the ACR is complete, current, credible, and useful for decision-making. A document may appear compliant while still failing to provide enough evidence to support award, acceptance, renewal, or continued operation. Conversely, a product may disclose accessibility limitations but provide strong evidence, a remediation plan, and contractually enforceable commitments that make the risk manageable.

Phase 1 ACR evaluation addresses this document-level problem. It helps determine whether the vendor's accessibility documentation is sufficiently credible to support the next procurement or governance decision. It is a screening, triage, and evidence-quality function; it is not a substitute for product testing.

2. EXPERT-DRIVEN MANUAL DESK AUDIT

In a manual desk audit, an accessibility expert reviews the vendor's ACR and prepares comments, findings, or a report. The expert may check whether the ACR uses the correct VPAT version, whether required sections are complete, whether the evaluation methods are credible, whether conformance claims are internally consistent, and whether remarks align with the reported support levels.

This model has important strengths. It can incorporate professional judgment, practical accessibility experience, and contextual interpretation. An expert reviewer may notice subtle contradictions that automated tools miss, and can assess whether a vendor explanation sounds plausible based on the product category or known accessibility patterns.

The model also has structural limitations. Manual review is relatively expensive and time-consuming. In the assumptions used for this report, a manual review costs approximately $1,500 per ACR report and can take up to two weeks to deliver. It is also difficult to scale when a procurement involves many vendors or when an agency needs portfolio-wide review across dozens or hundreds of systems. Finally, manual desk audits remain document reviews: they do not prove that the software is accessible unless paired with hands-on testing.

3. THE ACR EVALUATOR MODEL

The ACR Evaluator automates and standardizes many of the structured tasks involved in Phase 1 ACR review. It combines agentic AI, robotic process automation, expert-prepared rubrics, rules-based applicability logic, structured scoring, and report generation. These components are used to evaluate ACR completeness, credibility, internal consistency, applicability, evidence sufficiency, and risk indicators.

A central design feature is the automated applicability framework. Rather than relying only on a vendor's stated scope or a static checklist, the ACR Evaluator profiles a product across ICT Types and ICT Subtypes. That profile is resolved against the ICT Classification-to-Criteria Crosswalk Framework and Applicability Engine Specifications to identify applicable accessibility standards, requirements, evaluation criteria, and test methods.

The system therefore acts as a structured, automated first-line evaluator. It does not claim to verify actual product conformance by itself. Instead, it improves the consistency, speed, and cost-effectiveness of document-level ACR evaluation and helps determine when expert review or independent validation is warranted.

4. COMPARATIVE SUMMARY

The table below summarizes the two methods across the dimensions most relevant to a Phase 1 procurement decision.

Dimension	Expert-driven manual desk audit	ACR Evaluator
Primary method	Human expert reviews the ACR and prepares comments or a report.	AI-assisted and rules-based review using expert-prepared rubrics and automated workflows.
Typical commercial price	Approximately $1,500 per report.	$39 per report. Enterprise and volume discounts may apply for public-sector agencies.[2]
Typical turnaround	Up to two weeks.	Up to 10 minutes.
Scalability	Limited by expert availability and review capacity.	Designed for high-volume ACR intake, standardized review, and portfolio reporting.
Consistency	May vary by reviewer, workload, and interpretation.	Applies standardized rubrics, scoring rules, and repeatable report logic.
Applicability determination	Often relies on reviewer interpretation and vendor-provided scope.	Uses ICT Type / Subtype classification and crosswalk logic to determine applicable criteria.
Evidence assessment	Expert may comment on credibility and completeness.	ECS quantifies confidence in completeness, recency, independence, scope, and traceability.

Dimension	Expert-driven manual desk audit	ACR Evaluator
Human judgment	Central to every review.	Triggered when ECS is low or risk conditions warrant escalation.
Output	Narrative comments, annotated ACR, or desk audit report.	Structured report, scores, flags, risk drivers, recommendations, and audit trail.
Product validation	Not included unless separately scoped.	Not included in Phase 1 unless separately paired with IV&V or testing.
Best fit	Complex, ambiguous, high-risk cases requiring expert interpretation.	High-volume screening, standardized review, early procurement triage, and risk-based escalation.

Table 1 — Method comparison across the dimensions most relevant to a Phase 1 procurement decision.

Note: Price and timing assumptions reflect the commercial comparison used for this report. The ACR Evaluator is also offered with enterprise and volume discounts for public-sector agencies. Final pricing may depend on contract vehicle, agency volume, licensing model, and reseller terms.

5. COST AND TURNAROUND ANALYSIS

The economic difference between the two approaches is substantial under the assumptions used in this report. A manual expert-driven desk audit costs approximately $1,500 per report, while an ACR Evaluator report costs $39. On a per-report basis, the ACR Evaluator report costs about 2.6 percent of the manual desk audit price. Put differently, one $1,500 manual report is approximately equivalent in price to 38 ACR Evaluator reports.

Measure	Manual desk audit	ACR Evaluator	Interpretation
Approximate cost per report	$1,500	$39	ACR Evaluator is about 2.6% of the manual report price.
Approximate delivery time	Up to 2 weeks	Up to 10 minutes	Automated review supports earlier procurement triage.
Approximate reports for $1,500	1	About 38	Illustrative only; enterprise pricing may vary.

Table 2 — Cost and turnaround comparison under the stated commercial assumptions.

The time difference is also operationally significant. A two-week delivery window may be acceptable for high-risk, final-stage, or legally sensitive decisions. It can be less practical during market research, early vendor screening, or portfolio review. A 10-minute automated report allows procurement and accessibility teams to identify weak ACRs, missing evidence, and high-risk vendors earlier in the lifecycle.

The conclusion is not that every manual review should be replaced. Rather, the comparison supports a triage model. Automated review can screen broadly and consistently; expert review can then be focused on low-confidence, high-risk, ambiguous, mission-critical, or high-impact cases.

6. METHODOLOGICAL DIFFERENCES

6.1 Manual Interpretation versus Structured Rubric Application

Manual desk audits rely on the reviewer's expertise. This can be beneficial when the ACR includes unusual language or when the product context is complex. However, manual interpretation can also introduce variability. Two reviewers may assign different levels of concern to the same vendor statement. The ACR Evaluator applies expert-prepared rubrics consistently across documents. This improves repeatability and supports comparison across vendors.

6.2 Vendor Scope versus ICT Classification-to-Criteria Crosswalk

A traditional desk audit often starts with the vendor's stated scope. If the vendor marks a criterion as Not Applicable, the reviewer must determine whether that claim appears plausible. The ACR Evaluator starts from the selected ICT Type and ICT Subtype profile. The system resolves that profile against a crosswalk to determine which criteria should apply. This helps reduce two common scoping errors: false inclusion, where irrelevant criteria are applied, and false omission, where applicable criteria are left out.

6.3 Narrative Findings versus Quantified Risk Outputs

Manual desk audits often produce qualitative comments. These comments can be valuable, but they may be difficult to compare across vendors or aggregate across a portfolio. The ACR Evaluator converts document-level findings into structured outputs, including evidence confidence and risk indicators that can support procurement comparison, escalation thresholds, IV&V recommendations, and portfolio reporting.

7. EVIDENCE CONFIDENCE AND HUMAN-IN-THE-LOOP REVIEW

A key design feature of the ACR Evaluator is that automation does not operate without boundaries. The Evidence Confidence Score assesses whether the underlying ACR evidence is strong enough to support reliance on the automated report. ECS considers completeness, recency, version match, independence, scope coverage, and traceability.

When ECS is sufficiently low, the ACR Evaluator triggers expert intervention. This human-in-the-loop model helps prevent automated outputs from being treated as conclusive when the underlying evidence is weak. It also focuses expert time where it adds the most value: incomplete ACRs, questionable claims, outdated evidence, contradictory findings, high-risk products, and decisions requiring defensible judgment.

This design is important for public-sector use because procurement officials, IT teams, and Section 508 program managers need both efficiency and defensibility. The ACR Evaluator provides a structured first-pass review, while the low-ECS escalation path preserves expert oversight for cases where automation should not be the final decision point.

> **HUMAN-IN-THE-LOOP PRINCIPLE**
> Low confidence is not treated as neutral. If the ACR Evaluator cannot establish sufficient confidence in the evidence, the result should trigger expert review, additional documentation requests, independent validation, or other escalation before the decision is finalized.

8. STRENGTHS AND APPROPRIATE USE CASES

8.1 Strengths of Expert-Driven Manual Review

— Interpreting ambiguous or unusual vendor language.
— Assessing mission-critical, public-facing, or legally sensitive procurements.
— Evaluating equivalent facilitation claims or complex exceptions.
— Advising on remediation, contract negotiation, acceptance conditions, or risk acceptance.
— Conducting or supervising hands-on product testing and assistive technology validation.

8.2 Strengths of the ACR Evaluator

— Lower cost per report for broad Phase 1 screening.
— Faster turnaround for early procurement and portfolio triage.
— Standardized rubrics and scoring logic across vendors.
— ICT classification-based applicability determination.
— Evidence confidence scoring and risk-based escalation.
— Structured outputs suitable for procurement files, governance reporting, and audit trails.

9. SHARED LIMITATIONS

Both approaches share an important limitation: document review is not product testing. A well-structured ACR review can identify weaknesses in the vendor's documentation, but it cannot conclusively determine whether users with disabilities can complete real tasks in the product.

A vendor may claim support for keyboard access, screen reader compatibility, captions, or focus order, but those claims require testing to verify. Both manual review and ACR Evaluator Phase 1 review should therefore be understood as part of a broader accessibility governance workflow, not as a substitute for independent validation when the evidence or risk profile requires it.

10. RECOMMENDED OPERATING MODEL

The strongest model is not automation versus experts. It is automation plus experts. A practical operating model is:

— Use the ACR Evaluator to conduct fast, standardized Phase 1 ACR reviews.
— Use ECS and other risk indicators to determine confidence in the result.
— Automatically route low-confidence or high-risk reviews to expert intervention.
— Use expert reviewers for ambiguous, mission-critical, legally sensitive, or high-impact procurements.
— Use IV&V or product testing when documentation alone is insufficient.
— Maintain structured reports as part of the procurement and accessibility governance record.

CONCLUSION

Expert-driven manual ACR desk audits and the ACR Evaluator address the same fundamental problem: public-sector buyers need a defensible way to evaluate vendor accessibility documentation. Manual desk audits offer expert judgment and contextual interpretation, but they are relatively costly, slower to deliver, and difficult to scale. Automated ACR evaluation offers speed, standardization, and cost efficiency, but must be paired with human oversight when evidence is weak or risk is high.

The ACR Evaluator's principal contribution is not simply that it reduces review time or cost. Its more important contribution is that it transforms Phase 1 ACR review into a structured, repeatable, evidence-aware, and risk-based process. By combining agentic AI, robotic process automation, expert-prepared rubrics, applicability logic, evidence confidence scoring, and human-in-the-loop escalation, the ACR Evaluator provides a practical bridge between low-cost automated screening and high-value expert intervention.

For agencies evaluating multiple vendors, managing large ICT portfolios, or seeking more consistent accessibility procurement records, this hybrid model offers a defensible path forward: automate what is repeatable, measure confidence explicitly, and involve experts when the evidence or risk profile requires it.

REFERENCES

1. U.S. Access Board. Information and Communication Technology Standards and Guidelines — Revised Section 508 Standards and Section 255 Guidelines. 36 CFR Part 1194, 2017.
2. World Wide Web Consortium. Web Content Accessibility Guidelines (WCAG) 2.2. W3C Recommendation, 5 October 2023.
3. ETSI. EN 301 549 — Accessibility requirements for ICT products and services. V3.2.1, 2021.
4. Information Technology Industry Council. Voluntary Product Accessibility Template (VPAT) / Accessibility Conformance Report (ACR), Version 2.5.
5. Turner, M., Rathod, J. An Automated Framework for Determining ICT Accessibility Requirements. Revelo Software ACR Evaluator Technical Series, June 2026.
6. Revelo Software. Manual Document Review / Desk Audit Analysis. Internal working document, 2026.

AUTHORS & DISCLOSURE

Mark Turner, Ph.D.
Chief Product Designer, Revelo Software

Jatan Rathod
Analytics Engineer, Revelo Software

Disclosure. The authors are affiliated with Revelo Software, developer of the ACR Evaluator. The commercial price and turnaround assumptions compared in this report were specified for the analysis. The methodology and evidence base are disclosed in full to allow independent verification of the findings.