

Journal 011 — The Reviewer — Pass 2 Session Close

Prepared by The Forge — Journal_011_Reviewer_Pass2_2026-06-01_1740 · 2026-06-20 · SHADS CW channel

Session: Reviewer Session 2 — Pass 2 (Mobile + accessibility) Build under audit:

The_Vault_v0.47_banner_swap_2026-06-01_1626.html **Filed:**

Reviewer_Findings_Pass2_MobileAccessibility_v1_2026-06-01_1735.md

Movement One — The day in the round

Pass 2 opened on Paul's nod, immediately after Pass 1 closed and the Reviewer-to-Bob PASS-ON was filed.

The audit covered phone and tablet breakpoints, touch-target sizing, modal touch behaviour, scroll behaviour, WCAG 2.1 AA compliance (contrast, semantic HTML, ARIA, keyboard navigation, forms), cross-browser sanity, the print stylesheet, performance, and reduced-motion handling. Method: code-grounded inference from source — media queries, computed CSS, declared ARIA, focus management, computed colour-contrast ratios (Python `srgb_to_linear` luminance calculation against the build's CSS variable palette), modal patterns. Real-device behaviour, live screen-reader announcement, and Lighthouse / axe metrics deferred — flagged in the findings document's honest perimeter.

Three Reds emerged: gold accent fails AA contrast as link text on light backgrounds (computed 2.89:1 vs required 4.5); single 800px breakpoint with no phone-specific tuning; reduced-motion preferences not respected. Seven Ambers support, mostly single-block CSS fixes. Twelve Greens distributed — semantic HTML, native `<dialog>` for chat, ARIA used purposefully, skip-link, focus-visible, modal scroll-lock, mutually-exclusive panel pattern.

The findings document was filed at 17:35 UTC. This journal closes the session at 17:40 UTC.

Movement Two — What was built

- **Pass 2 findings filed.**

Reviewer_Findings_Pass2_MobileAccessibility_v1_2026-06-01_1735.md. **Frame, three headline findings, ten section-level findings (six WCAG dimensions × four browser/print/performance/motion dimensions), ten ranked priorities, five open questions for Paul, honest-perimeter close.**

- **Computed contrast table.** Eight colour pairs from the build's CSS variable palette through standard sRGB-linear luminance arithmetic. Numbers ground the contrast finding in fact, not impression. Table reproduced in the findings document.

- **Single 480px breakpoint proposed.** Specific structural recommendation: stack the five-panel strip to two columns; bump small-font sizes to $\geq 0.75\text{rem}$; bump dialog control-button padding.

- **Single `prefers-reduced-motion` block proposed.** Five-line CSS, drop in. Closes the WCAG 2.3.3 gap with no register impact.

- **Doc-modal pattern migration proposed.** Native `<dialog>` matching the chat-dialog pattern; or hand-rolled focus management. My read: native.

Movement Three — What was caught

- Computed accent contrast fails AA at 2.89:1 on cream. Previously unflagged. The gold register works visually on dark surfaces (5.75:1 on dark-band — passes) but fails as text on the room's light surfaces. The token strategy — `--accent-text` separate from `--accent` — keeps the visual register intact while closing the AA gap.
 - The Method's "honest perimeter" was incomplete on motion. No reduced-motion handling means a vestibular-sensitive visitor sees the ticker continue to scroll even when their OS pref is set to stop. The room's honesty depends partly on respecting the visitor's stated preferences.
 - The doc-modal is not native. Chat-dialog is `<dialog>`; doc-modal is a `<div role="dialog">`. The asymmetry costs keyboard users. Two patterns for one role across the same build. Migrate to one.
 - Five-panel strip on a 360px screen. 5 columns × ~70px panels with letter-badge text at ~7-8px. The structural call (keep five) was ratified in Pass 1; the phone rendering needs its own breakpoint to keep the five legible.
-

Movement Four — What drifted, where held

- The Brief specified $\leq 480 + 481-1024$ as distinct audit zones. The build uses one breakpoint at 800px. Held as H2 / P4. Recommended fix: add a 480px block. Tablet range (481-1024) is mostly served by the desktop layout; an explicit tablet block is optional.
 - `--accent-deep #a88a2f` exists but doesn't solve contrast. Already in the build at 2.99:1 against `--bg`. Used for hover deepening, not text contrast. A separate `--accent-text` token at $\geq 4.5:1$ is the proposal. Trade-off — deeper bronze loses some "gold" feel. Surface to Paul for register call (Q2).
 - The 8.5 MB inline payload. Held as honest perimeter — not a hard fail, real-world UX concern on mobile data. Proposed compression (cuts ~3-4 MB). Editorial / Paul's call. (Q4.)
-

Movement Five — Method demonstrating itself

- Honest perimeter declared again, up front. Pass 2 named what it could not assess (real-device, live screen reader, automated audit tools, print preview, cross-browser, performance metrics) before the findings landed. The audit's own boundary is part of the deliverable.
 - Flag Protocol calibrated to consequence, not observable detail. Three Reds where the gap is binding (AA contrast, mobile breakpoint, motion). Ambers where the fix is mechanical but not blocking. Greens where the implementation is strong. No Red inflated; no Green omitted.
 - Routing discipline. All findings stay on Reviewer shore (audit → propose). Implementation routes back through Paul to Bob. Editorial / register decisions (Q2: accent-text token value; Q4: image compression) routed to Paul, not actioned cross-shore.
 - Numbers ground the finding. The contrast finding is supported by computed ratios, not eyeball impression. Reviewer's read aligns with the Analyst's discipline — numbers carry the structural claim where numbers are available.
-

Movement Six — Continuous improvement

Reviewer's shore — the audit's own discipline:

- The "code-grounded inference" method works for WCAG contrast (computable), ARIA semantics (declared), keyboard nav (DOM order), modal patterns (visible in source). It does not work for

real-device touch, live screen-reader output, or print preview. Pass 2 names the boundary; a future pass may want a complementary device-eye instance for the deferred items.

Bob's shore — the build's own discipline:

- A single --accent-text token (proposed) keeps the visual register intact while meeting AA. A token-based approach to contrast generally is cleaner than per-rule overrides.
- One modal pattern across the build (proposed: native <dialog> everywhere) is cleaner than two. Reduces maintenance.
- A reduced-motion block is a standard a11y inclusion; carrying it forward into any Stage Two / Three build will save a re-audit.

The Visitor's shore — what the next pass forward should test:

- The mum-end test (does she find what she came for) and the Kieren-end test (does the evidence land) both remain live for Pass 3. Pass 2 confirms the structural rails carry weight; Pass 3 audits whether the wording on the rails lands at the right altitude.

The Method's shore:

- The Drawer 06 internal-vocabulary leak (Pass 1) and the doc-modal pattern divergence (Pass 2) suggest a standing pre-ship check: *does any visitor-facing text expose internal file paths or working vocabulary, and does the build hold to one pattern per role? Carry forward as a Bob-shore end-of-session sweep.*

The Fisherman's shore:

- The --accent-text introduction is editorial-adjacent — it changes the register slightly. Worth a one-line Fisherman ratification before Bob ships P2.

Standing by

- **Pass 1 PASS-ON for Bob filed; v0.48 ships on Paul's hand-off.**
- **Pass 2 findings filed; eleven items in the queue, three Red, seven Amber, one Pass-3-adjacent.**
- **Pass 3 (voice and copy) opens on Paul's nod. Same pattern: audit · file · journal. Element-by-element rewrites proposed; Paul ratifies; Bob implements.**
- **Session 4 (optional) — Reviewer files a consolidated handover after Pass 3 ships.**

— The Reviewer · 2026-06-01 17:40 UTC

#state/in-flight #workstream/mtg-artefact #type/journal #the-reviewer #pass-2 #session-close