HealthCode Analysis — Comprehensive Technical Overview

Native Elementor publishing, verified security boundaries, US/EU activation requirements and a staged scale architecture

A public reference for clients and developers. This edition reflects the deployed native WordPress release and its recorded verification. It separates working features from future architecture and business decisions.

Developed by Zahidul Islam. Edition verified: Sep 14, 2026 · Documentation updated: 24 September 2026

Open the live site · Source repository

Latest status — reviewed 14 October 2026

Review scope. This review compared the existing Google Doc with the current project overview, changelog, roadmap, threat model and test records. The September additions are already present. This summary brings the latest recorded changes and open work forward. This is a documentation review; it does not renew the earlier deployment, legal-research or test evidence.

Current reachability — 14 October 2026. One header-only HTTPS request to the public site returned HTTP 200. It also showed the duplicate Strict-Transport-Security header already tracked in the roadmap. This confirms a responding endpoint at the time of the request; it does not establish current browser behavior, authenticated editing, sustained availability or deployment parity.

Latest recorded update — 24 September 2026. The repository gained architecture, operations, incident-response, disaster-recovery, development, threat-model and legal/compliance documentation. That review recorded 316 passing offline checks, a failing retained AskMe test and two runtime-dependent integration suites that were not run. No new test run is claimed on 14 October.

Recorded native release — 14 September 2026. The WordPress and Elementor Free release has 66 native layouts with zero HTML widgets, editable header and footer, browser-only library/search and six educational tools. Its recorded evidence includes 13 browser scenarios, 11 motion/accessibility states, 21 HTTP/security checks, MFA/editor access, a separate-runtime restore and 141-file deployment parity. These remain dated release results rather than proof of today's live state.

Open quality and security work. The September roadmap still records a broken AskMe test harness, eight legacy lint findings, guarded CI without an executed run, action tags awaiting SHA pinning, machine-specific paths, trusted-proxy configuration ownership and local Docker port binding. The legacy snapshot's public AskMe route is separately recorded as an exposure: it is not used by the native library, and decommissioning or origin restrictions plus rate limiting remain open decisions.

Next milestones and acceptance. Owner visual acceptance, screen-reader and physical-device review, external monitoring and timed recurring restore drills remain unconfirmed in the reviewed records. Observability precedes edge delivery, redundant origins/database and staged capacity qualification. Ten thousand to one million concurrent readers and 99% availability remain targets. Collection, monetization and optional AI integrations remain disabled pending their individual activation gates.

Historical material retained. All 29 existing sections remain below, including earlier results, source links and limitations. Legal summaries and software-version references retain their original research context; this review does not certify them as current. In particular, the Bangladesh law summary remains explicitly unverified against the official gazette. Refresh applicable primary sources before a commercial launch or changed data flow.

1. Reading guide and evidence policy

Clients can start with the product, design, editable-content, legal, capacity and acceptance sections. Developers should also read configuration ownership, delivery, recovery, security and verification. The repository contains a current native implementation alongside retained automation and the original design; those components have different operating roles.

Implemented means a capability exists in the inspected code. Verified means a named check exercised it in the stated environment. Planned means additional engineering or infrastructure remains. A target is not a measured result. No part of this document promises an unhackable system, universal legal compliance, guaranteed search placement or proven million-user capacity.

The 24 September 2026 documentation update adds a public engineering and compliance set to the repository: architecture, roadmap, threat model, incident response, backup and disaster recovery, operations runbook, development guide and a legal pack. That update covered documentation, repository settings and offline re-verification only; it made no production server, routing or content change. The release evidence in this overview remains the recorded 14 September verification unless a paragraph states a later check.

2. Executive overview and current release

HealthCode Analysis is a medical-technology editorial build: articles, reviews, a searchable library, a device-local reading list and six educational browser tools. Its visual direction combines a dark graphite foundation, mint accents, copper illumination, large editorial typography and an illustrated neural brain. The same design is editable through native Elementor Free layouts in WordPress.

The public HTTPS release passed 13 browser scenarios, 11 motion/accessibility states and 21 HTTP/security checks. MFA-protected administration and the Elementor editor opened successfully. All 141 versioned release files and the project routing configuration matched the local release by SHA256. A separate 28-request test confirmed that forged application client-address headers did not bypass the login limiter.

The site is an independent build, not an operating healthcare provider or validated clinical product. Imported sample articles remain unverified and excluded from indexing until editorial approval. Ads, affiliates, analytics, newsletter subscriptions, public accounts and server-side health-data collection are disabled. No new paid service or runtime paid AI inference is required by this release.

3. Product purpose and boundaries

The project shows professional publication design, practical native CMS authoring, browser interaction, controlled deployment and documented engineering boundaries. It preserves 63 original public routes, Unicode navigation and a real missing-page response. Additional policy and CMS records bring the restored database to 68 published pages; route counts, generated-layout counts and database counts measure different things.

It does not diagnose, prescribe, validate medical images, claim regulatory clearance or provide emergency care. Search retrieves library matches rather than generating medical answers. Calculators and worksheets are educational tools; their presence is not evidence of clinical validation or a medical-device exemption.

4. Frontend design system and motion

The retained Neural Frontier composition uses a two-line publication headline beside the brain artwork, with a centered wide-screen frame that gives the right side breathing room. Typography, gutters, card

rhythm, outlined and filled buttons, directional icons, review stacks and the lighter toolbox section work together as an editorial medical-technology publication.

The hero sequence occupies 200svh. Compared with the previous 175svh sequence, available scroll travel increased from 75vh to 100vh, reducing frame progress per unit of scroll by 25%. The user can pause and resume motion. Reduced-motion preferences, narrow screens and failed media loading display the still illustration; the mobile test verified that the film was not requested.

The artwork is explicitly disclosed as AI-generated illustration, not a medical scan. One image and one eight-second film were created under a previously approved visual pilot and optimized locally with FFmpeg. A Blender composition draft and the original frontend are retained. There is no per-visitor image-generation or video-generation API call.

Desktop, laptop and mobile comparisons included the owner's full-HD viewport, 1440-pixel desktop and 390/320-pixel mobile widths. A theme-added mobile header inset and an overflowing disclosure were corrected. Visual comparison is evidence of the implementation; the owner still decides whether the result matches the intended design closely enough.

5. Native Elementor editing and ownership

The generated migration contains 66 native page layouts using Elementor containers and free heading, text, image, button and video widgets. The build uses zero HTML widgets. Header and footer templates use the free Ultimate Addons/header-footer builder. Elementor Pro is not required.

Custom functionality remains in a small scoped plugin: library controls and educational-tool shortcodes supply behavior that standard layout widgets do not implement. The conversion contains 28 library-filter widgets and six tool shortcodes. This is an explicit functional exception, not a whole-page HTML wrapper or iframe substitution.

Representative heading and button changes were saved through visible Elementor controls; a representative image change used Elementor's native settings command. Public HTTP responses showed the edits, originals were restored, and custom behavior attributes survived editor reload. The restored server and public site also passed authenticated editor access checks without modifying content.

ACF provides search description, same-site social-image URL and an editorial indexing switch. Social metadata accepts both imported relative paths and valid same-origin absolute URLs; invalid or external values fall back to the configured illustration. Five actual field-save/render cases passed with the original value restored.

6. Architecture and state boundaries

Public reader → existing CDN/TLS proxy → project gateway → anonymous HTML cache → WordPress on cache miss → isolated database. Static media is served from the site, and library search and educational tools run in the browser. Authenticated editors use the uncached WordPress administration path.

WordPress owns editable content after import. Source code and maintained configuration own application behavior and release construction. The original frontend is a preserved reference and rollback artifact. Importing generated layouts is an explicit replacement operation; maintainers must reconcile any newer editor changes before doing it again.

The runtime has separate web, database and gateway services, persistent database/content volumes and a tools-profile CLI. The database network is internal. The gateway is nonroot, read-only, capability-dropped and bound to a loopback host interface behind the TLS proxy. Per-service CPU and memory bounds limit resource consumption on the shared host; they do not create horizontal scaling.

The old static asset route is retained so already-open reference pages can still retrieve their CSS, scripts and film after migration. New native assets use the scoped plugin path. Only the project's hostname routing was changed; other hosted projects were outside the deployment scope.

7. Configuration contract and developer ownership

Environment-dependent settings enter through the typed NativeEnvironment boundary and documented environment examples: image references, origins, port, database settings and administrator setup. Actual credentials remain in ignored private recovery material. Public documentation contains no passwords, tokens, account emails, server IPs or private operating paths.

Maintained JSON owns publication/privacy features, editorial metadata fields, public notices, gateway delivery limits, icon palette, tool content and formulas, capacity assumptions and local smoke-test bounds. The restore policy separately owns import-only SQL compatibility. The original frontend has its own site and media configuration boundary.

A developer should be able to change supported deployment settings through environment/configuration data and change editorial copy through maintained content or WordPress. Protocol identifiers such as the SVG namespace and schema.org vocabulary remain fixed constants. Runtime provider/model identifiers must not be embedded in business logic; the current public runtime does not call a model provider.

The local importer validates loopback site origins and rejects remote Docker endpoints from environment variables or the active context before any Compose mutation. It pins a validated local IPC endpoint and removes inherited Docker selection/TLS overrides. A local-looking URL alone is not treated as proof of a local daemon.

8. Delivery lifecycle and safe mutation

The release process starts with acceptance criteria and a current-state snapshot. Files that will change are compared with the live versions before local edits. Newer production edits must be reconciled locally first. Build and verification then produce a reviewable release, followed by isolated restoration, public activation and final file parity.

Before native production restoration, the deployment compared editable page content, Elementor layouts, page settings and ACF publication metadata against the authoring copy. It saved a database backup before restoring the prepared production-origin export. The origin was checked for expected native content, correct canonical identity and production environment before routing changed.

The project routing candidate was rendered locally, validated, compared against the prior live configuration and installed only for the intended hostname. Activation failures restore the retained static configuration. First-run checks are inside the rollback-protected operation; repeat checks read the current routing and public content rather than trusting an old success checkpoint.

9. Library, search and reading list

Library search and Ask the library download a public content index and perform matching in the browser. They do not call the retained AskMe Worker or a remote generative model in this release. Empty results and unsupported queries receive a bounded navigation response rather than invented medical guidance.

Saving a story stores its identifier in browser storage on that device. The interface supports persistence, removal and an empty state. If storage is blocked, the user receives actionable feedback. The reading-list route is private/no-store at the response layer even though the server does not hold a personal reading profile.

The browser audit exercised search results, empty states, Escape handling and focus return, saving/removing stories, archive category filtering, mobile navigation and blocked storage. It found no JavaScript errors or external network requests in the exercised public flows. This observation does not mean hosting/CDN connection metadata or administrator data does not exist.

10. Six educational tools

The BMI calculator shows adult screening calculations with context and input bounds. The eGFR calculator shows an educational estimation flow. The AF score worksheet supports a bounded educational scoring exercise. None is presented as a diagnosis or treatment recommendation.

The PICO research worksheet helps structure a research question. The Medical prompt studio assembles a prompt locally; it does not submit that prompt to a model. The Image workbench processes a selected image in the browser and supports an output download. It must not be described as clinical interpretation, diagnostic image validation or guaranteed anonymisation.

All six real browser flows passed, including the image output/download path. Form values are not intentionally sent to the server or embedded in query URLs. Future telemetry, uploads or model integrations would change this data-flow boundary and require a new security and legal review.

11. Search-engine and search-agent optimization

The native publication layer generates page descriptions, canonical URLs, Open Graph/Twitter image metadata and truthful structured data. Eligible pages appear in WordPress sitemaps. Administrative, author-enumeration and private reading-list surfaces do not become normal discoverable publication pages. Query variants canonicalize to the base page URL; unverified sample material has an explicit noindex boundary.

The ACF indexing field is an editorial gate: verify claims, source rights and page purpose before enabling it. Local and staging environments remain noindex regardless of the field. The production homepage's indexing eligibility was checked over HTTP. Unverified imported stories are not disguised as reviewed medical articles, ratings or expert endorsements.

Search-agent work follows ordinary crawlability, semantic structure, descriptive links, accessible controls and reliable page facts. Google does not require a special AI markup file, and technical eligibility does not guarantee indexing or an AI citation. OpenAI search discovery and training crawlers have different controls; the configured policy disallows GPTBot while retaining appropriate public discovery. Robots rules are cooperative instructions, not authentication or access control.

Search Console ownership, search rankings and agent citation rates have not been measured or claimed. A client must complete actual account verification and editorial publication work before treating search visibility as a delivered commercial outcome. Google AI search guidance and OpenAI publisher guidance explain the distinction.

12. Runtime versions and maintenance

The deployed inventory reports WordPress 7.1, PHP 8.5.10, Elementor 4.2.4, Ultimate Addons 2.9.4, ACF 6.8.10 and Two Factor 0.16.0. The Hello Elementor theme is 3.5.1. Active plugins report no pending update in the checked inventory. These are release-time observations, not a promise that future versions will never appear.

Official package inventories and checksums informed upgrades. The core checksum check passed with an expected extra Docker configuration-template warning; the gateway denies that template's public path. Two Factor's compatibility metadata lagged the installed WordPress version, so actual password/code authentication and editor access were tested rather than assumed from metadata.

Minor core automatic updates are enabled; major changes and plugin updates still require backups, compatibility checks, actual editor/public-flow tests and deployment parity. No paid or nulled plugin is part of the native build. Inactive bundled plugins were retained rather than removed without need.

13. Authentication, authorization and attack-surface controls

The deployed administrator account requires a password and TOTP second factor. Enroll and verify MFA separately for any future privileged account. Invalid local codes were rejected; valid local, restored-server and public logins opened the native editor. Recovery material is stored privately. A login page's existence is not proof that a second factor is enforced, so the password-only boundary was exercised.

The native runtime disables application passwords, XML-RPC, comments/pings and unused public author enumeration. Unauthenticated REST user listing is rejected while the APIs needed by authenticated Elementor remain usable. Dashboard source-file editing is disabled. These choices apply to the current native runtime; the retained legacy automation bridge has a different authentication model and is not implicitly enabled.

The gateway denies configuration, hidden and backup-like paths, directory listings and execution of uploaded PHP-like files. The database is not publicly exposed. Login requests have a narrow rate limit; ordinary publication reads are not subjected to that login threshold.

The proxy trusts a specifically configured host gateway, while the outer routing overwrites the application client-IP header. Official CDN address ranges constrain when the provider's connecting-address header is accepted. A 28-request production test used forged application headers: 11 requests succeeded and 17 received HTTP429, with the actual client identity confirmed privately in limiter logs. This is a bounded trust-boundary test, not a denial-of-service resistance benchmark.

A published STRIDE threat model records 17 threats with their controls and status. It marks partial controls openly, including single-origin denial-of-service exposure, the absence of central audit-log retention and CI actions pinned by tag rather than commit SHA. Vulnerabilities are reported privately through the repository's security policy, which defines scope and safe harbor, or through GitHub private vulnerability reporting, which is enabled.

14. Browser policy, caching and privacy isolation

Public responses include a content security policy, MIME-sniffing protection, same-origin framing policy, restricted browser permissions, referrer policy and HTTPS transport policy. Registered WordPress inline scripts receive hashes so public HTML can remain cacheable without allowing arbitrary inline scripts. Inline style compatibility remains a conscious Elementor constraint; this is not a claim of perfect browser isolation.

Private/no-store is the default. Only eligible anonymous successful page responses become cacheable. Requests with cookies, authorization, query strings, private routes or non-read methods bypass the shared HTML cache. Responses that set cookies are not stored. Public and private response headers were checked as well as cache HIT/BYPASS behavior.

The gateway uses bounded cache storage, cache locking and controlled stale/background updates. Public cache lifetime is configured at 30 seconds with a further 30-second stale-while-revalidate window. After importing layouts, generated Elementor CSS and gateway cache must be invalidated together to avoid old markup referring to new styles.

Cloudflare proxying does not itself prove HTML edge caching: HTML is not cached by default. The implemented reverse-proxy cache is an origin-level scaling boundary. CDN HTML caching, distributed invalidation and large media delivery remain separate work. Cloudflare cache guidance

15. Privacy and current data inventory

Public reading involves the browser, CDN and hosting service, which can process connection information. Administration uses authentication cookies, account details, a database and backups. Therefore “no analytics” must not be rewritten as “no personal data exists.” Security logs and infrastructure metadata require defined purposes, access and retention.

Library queries, educational inputs and image processing stay in the browser in the exercised flows. Saved story identifiers persist only after a user saves them and can be removed. There is no configured public account, newsletter, advertising, affiliate tracking or analytics collection flow.

The legal operator, establishment and working privacy contact have not been supplied. The public site must not invent them. Before accepting real public submissions or operating commercially, the client must identify the controller/operator, service providers, lawful purposes, retention and applicable request-handling procedures. An independent-build label is not a legal exemption.

16. European Union: applicability and required activation work

GDPR territorial scope depends on establishment and relevant offering/monitoring activities; simple global availability is not the entire test. A future plan to target EU users requires an operator-specific assessment. Where applicable, identify the controller and contact, document purposes and lawful bases, give transparent notices, minimize collection and support relevant access, correction, deletion and other rights. EDPB territorial scope and small-business guide

Security duties concern confidentiality, integrity and availability. Health-data collection or profiling can change the need for risk assessment, a DPIA, processor contracts and stronger access controls. International transfers require their own assessment and an applicable mechanism; server location alone does not settle them. EDPB security guidance and Commission transfer rules

For a relevant personal-data breach, assess the GDPR regulator-notification framework promptly, including the 72-hour rule where notification is required, and separately assess communication to affected people. Record the decision rather than assuming every incident has identical duties. This is distinct from US health-breach rules. EDPB breach guidance

Cookies and similar storage need a purpose-specific assessment under applicable ePrivacy rules. The absence of advertising does not remove every storage question. Necessary authentication and a user-requested saved list differ from optional tracking. Before adding nonessential analytics or advertising, implement any required prior consent, equally usable refusal, withdrawal and actual pre-consent blocking; a decorative banner is insufficient. EU online privacy guidance

The reviewed Commission guidance states that AI Act Article 50 transparency duties apply from Aug 2, 2026. Provider and deployer duties differ. The current search is retrieval, not a generative chatbot; AI-generated brain artwork is disclosed. A future model integration needs a role, transparency, data-transfer and intended-use assessment. Commission AI transparency guidance

Medical-device software qualification depends on intended purpose and functionality, including claims and the information produced. A footer disclaimer and browser-only calculation do not establish an exemption. Before patient-specific clinical use, obtain an appropriate qualification/classification decision and address required quality, evidence, risk and conformity work. MDCG software guidance

The European Accessibility Act covers specified products and services, not every informational blog automatically. E-commerce and other covered activities can change applicability; entity-specific exemptions require facts. Future sales also require consumer-information, cancellation, pricing, tax and country-specific review before checkout activation. EU accessibility guidance

17. United States: health, privacy and publication

HIPAA applies to covered entities and business associates, not every health-themed site. A future client relationship involving protected health information can change the result. TLS, a WordPress plugin or a privacy page does not establish HIPAA compliance. HHS covered-entity guidance

The FTC Health Breach Notification Rule can cover specified non-HIPAA personal-health-record vendors, related entities and service providers. Its health-app scope and unauthorized-disclosure rules must be assessed before connecting data sources or retaining identifiable health records. The current tools have not been determined to be a covered PHR service. FTC health-breach guidance

State consumer-health laws can extend beyond HIPAA. Washington's My Health My Data framework is relevant before collecting or inferring identifiable health interests. Its requirements should not be generalized to every state. California CCPA applicability depends on actual business, threshold and processing facts; where applicable, rights, notices, sale/sharing opt-outs and recognized preference signals must work. Washington health-data guidance and California CCPA guidance

COPPA can apply to child-directed services and services with actual knowledge of collecting personal information from children under 13. The site does not add children's accounts, age collection or targeted advertising. Such features require a separate assessment before activation. FTC COPPA guidance

Material affiliate or endorsement relationships require clear, conspicuous contextual disclosure. A generic footer does not replace an appropriate nearby disclosure. False reviews, fabricated testing experience, invented credentials and unsupported ratings must not be used to improve conversion or schema completeness. The imported stories are sample content, not verified product testing. FTC endorsement guidance and review-rule guidance

FDA clinical-decision-support guidance is relevant before commercial patient-specific decision support. The tools have not received an exemption or clearance determination. Accessibility duties under the ADA also depend on entity and context; technical keyboard checks do not settle every legal question. FDA clinical decision support and DOJ web accessibility

This assessment uses primary regulator guidance and the cited technical sources. Direct retrieval of the GDPR legislative text was blocked by browser verification; Washington guidance was available through official indexed material while direct retrieval returned403. The report does not claim a complete statutory review of every US state or EU Member State. Real launch requires current jurisdiction-specific review.

The public compliance register adds further US rows. COPPA is not applicable while the site remains general-audience with no accounts; reassess before accounts or age-specific content. CAN-SPAM applies once a commercial newsletter exists: sender identification, a physical address and unsubscribe requests honored within 10 business days. DMCA §512 safe harbor requires a designated agent registered with the US Copyright Office before user uploads are hosted; a copyright-takedown process is already documented.

Other jurisdictions are recorded in the same register. UK GDPR and the Data Protection Act 2018 impose GDPR-equivalent duties, an ICO fee/registration assessment and UK transfer mechanisms where the operator is UK-established or targets people in the UK. Bangladesh's Personal Data Protection Act 2026, reported to replace the 2025 Ordinance, is relevant if the operator is established there; secondary sources report phased implementation through 2027, the text has not been verified against the official gazette, and local counsel review is required.

18. Commercial feature activation gates

Real publication requires claim verification, source/media rights, truthful authorship, conflict disclosure and an editorial approval record. Medical authority must not be invented to fill a template. Eligible-only indexing is a working boundary, not a replacement for editorial review.

Analytics and advertising require a data-flow map, applicability/consent decision, tested refusal and withdrawal, truthful notices and provider review. Affiliate links require an actual relationship and contextual disclosure. Newsletters require a real operator/contact, an appropriate permission basis, unsubscribe, suppression and retention handling.

Public accounts require authorization tests, secure recovery, abuse controls, deletion/export procedures and retention. Health information requires regulatory qualification, an applicable processing basis, contracts, risk/impact assessment, access controls and a breach procedure. Generative AI requires role/disclosure analysis, provider configuration and safeguards against unintended health-data transfer or training use.

These features remain disabled. No placeholder policy should be presented to a client as proof that a future business is already legally operational. No new paid integration is activated by this roadmap.

A repository legal pack supports these gates: a medical disclaimer, a verified cookie and browser-storage inventory, a service-provider register, an AI-transparency statement, an accessibility statement, an editorial and copyright policy, and launch templates for a privacy policy and terms of use. The templates are not in force, the operator fields remain to be completed, and none of the pack is legal advice. The site notices continue to be maintained separately.

19. Accessibility and resilient interaction

The implementation includes keyboard navigation, visible focus, search-dialog focus restoration, a skip link, responsive navigation, readable fallback states and reduced-motion handling. Hidden duplicate image links are removed from the tab order. Long imported URLs wrap within narrow article columns rather than being concealed with overflow clipping.

Actual browser tests exercised these behaviors, including stacked review links and mobile navigation. They are useful engineering evidence, not a full screen-reader, physical-device, cross-browser or WCAG conformance assessment. A client-facing accessibility statement must distinguish tested features from known limitations and provide a real contact once the operator is established.

A published accessibility statement now targets WCAG 2.2 Level AA and makes no conformance claim. It lists known limitations, including no recorded screen-reader or physical-device testing and unaudited alternative text and contrast in imported articles, and sets an audit roadmap: automated checks across all 63 routes, manual keyboard and screen-reader passes on key journeys, a contrast audit and published remediation results.

20. Capacity target: 10k–1M concurrent readers

The codebase establishes practical scaling boundaries: browser-side tools/search, anonymous HTML caching, explicit private bypass, cache locking, bounded services, separate authoring and publication concerns, a configuration-owned capacity model and safe local test tooling. These choices reduce unnecessary origin work and make a larger architecture feasible to develop and test.

The current single-server site has not proven 10,000, 100,000 or one million simultaneous users. Concurrent readers differ from simultaneous requests, requests per second, authenticated database sessions and live video streams. A credible claim must define geography, browsing interval, page mix, cache hit ratio, media bytes, cold starts, authentication and arrival pattern.

The example model assumes a 30-second page interval, eight requests per page, 99% edge hit ratio and 500 kB per page view. At one million readers, those assumptions imply about 266,667 edge requests/second, 2,667 origin requests/second and 133.3 Gbit/second delivered at the edge. These are calculations, not observed capacity. The 500 kB example excludes a fresh multi-megabyte hero-film download; video traffic needs a separate budget.

A free feature or proxied domain does not promise free delivery at that traffic level. Provider limits, bandwidth terms, quotas, test-generator capacity and operating costs require explicit review before any large test or additional infrastructure. No high-load test was aimed at the shared production server.

21. Staged scale architecture and qualification

Stage 1 establishes the measured baseline: correct native editing, secure private/public boundaries, restore evidence, cache behavior, CPU/memory/database observations and page weight. The implemented local smoke tool sends only 40 requests with four workers after warmup, refuses nonloopback targets, disables inherited proxies and does not follow redirects.

Stage 2 scales anonymous reading through a validated CDN HTML strategy or a static publication pipeline generated from native WordPress. Elementor remains the authoring system. Version assets, test invalidation and cache poisoning/query variation, separate media delivery from HTML and protect the uncached editor. CDN HTML caching is planned; the origin gateway cache is implemented.

Stage 3 introduces redundant origins, health-aware routing, shared upload/state design and a database replication/recovery strategy where dynamic responses remain necessary. Merely copying containers does not replicate state. Test origin, database and relevant regional failures. Static publication can keep reading available during authoring outages, but freshness becomes a separate objective.

Stage 4 qualifies each target with authorized, production-representative gradual tests: cold/warm caches, purges, bursts, sustained traffic, media, bots and failure under load. Record release/configuration identity, duration, geography, request mix, errors, latency percentiles, resource headroom and dropped generator iterations. A test script or architecture diagram is not a passing load result. k6 threshold guidance

The completed bounded local smoke returned 40/40 successful cache hits at four workers, approximately 67.91 requests/second, median 6.995 ms and p95 around 9.75 ms. It measured loopback HTML only and loaded no browser assets. It must not be extrapolated into a production concurrency claim.

22. Availability target: 99% and recovery objectives

The proposed public availability SLI is successful external automated page probes divided by all scheduled eligible probes over a rolling 30-day window. A success requires valid TLS, expected status and content within a documented deadline. Monitoring gaps are reported separately; missing probes are not successes. Administrative availability and publication freshness are separate indicators.

The 99% target is an SLO, not a contractual SLA or measured achievement. As continuous-time arithmetic, 1% of 30 days is 432 minutes, or 7 hours 12 minutes. Sampled probes estimate availability and do not establish exact outage duration. Independent locations, cadence, deadline, retry treatment and incident ownership must be configured before the observation window starts.

External monitoring, automated failover and long-term availability evidence are planned. Backups establish recoverability, not instantaneous availability. Recovery-point and recovery-time commitments should follow repeated representative restore exercises. Maintenance and outages must be reported honestly; agree any contractual exclusions explicitly. Google SRE guidance

A proposed error-budget policy governs releases within each 30-day window. Below 50% of the budget consumed, releases proceed normally; from 50% to 75%, each release needs a reviewed rollback plan; from 75% to 100%, only reliability and security fixes ship; beyond 100%, a feature freeze applies until a post-incident review is complete, and the miss is reported publicly. Details are in the scalability and reliability plan.

An incident response plan defines four severity levels, with SEV-1 incidents (site down, tampered content or a suspected personal-data breach) acknowledged within one hour. It includes a

personal-data breach decision tree and a post-incident review template. The plan is documented, not yet exercised under a live monitoring window.

23. Backup, restoration and rollback evidence

A fresh local SQL/content snapshot was taken with authoring writes stopped and then restored on a separate isolated server runtime. The restored database contained 68 published pages and 156 valid Elementor layouts. Public interactions, film behavior, MFA and the native editor worked on that restored runtime before production activation.

Database readiness requires an authenticated TCP query against the configured database. A process ping was insufficient because the image briefly runs a socket-only initialization server. Transformed WordPress exports also required connection-scoped SQL compatibility for older zero-date defaults; global database modes were not relaxed. The exact failed import was reproduced and the compatible restore verified separately.

The production transition saved another database backup and compared editable facts before applying the prepared production-origin export. The restore policy is included in the 141-file recovery manifest. The original static release remains available for routing rollback; recovering WordPress content and rolling back public routing are distinct operations.

The backup and disaster-recovery plan proposes recovery objectives: a broken deploy restored within 30 minutes with no data loss by release rollback; database corruption within 4 hours and host loss within 24 hours, each from a backup no older than 24 hours; and a CDN/DNS misconfiguration within 1 hour. These are proposals. They become commitments only after two consecutive timed drills meet them. Scheduled automated backups and off-site retention enforcement are planned.

24. Verification, security scanning and limitations

The native behavior/configuration suite passed 11 tests; scoped Ruff lint and Mypy across six native Python files passed. The blocking Bandit medium/high gate passed. General Semgrep rules covered 22 applicable rules over nine files; the PHP pack covered 23 rules over four PHP files, with no findings in those runs. Static analysis reduces risk but does not prove that every vulnerability is absent.

The final pre-commit candidate source scan covered 178 files, approximately 6.29 MB, with no Gitleaks findings at that point; the final staged diff is scanned again before commit. The earlier 43-commit history scan was clean. Private credentials, environments, recovery records and master dossier remain ignored. Public cryptographic/protocol identifiers are distinguished from actual secrets.

The original frontend suite passed seven tests. Retained offline automation groups passed 71, 46 and 38 checks. Blind discovery of an old integration script failed because it assumes its original Docker runtime and executes at import time; the file was preserved and that exception is not counted as a native pass. No fabricated aggregate coverage percentage or CI success badge is used.

The current security workflow is manual and guarded, with no executed CI run claimed and no paid AI review job enabled. Real HTTP, browser and editor checks complement code tests. Fresh-context review corrected deployment, importer, social metadata and workflow defects before acceptance and found no outstanding material source blocker in the final reviewed change.

A 24 September 2026 offline rerun passed 316 checks: 11 native behavior/configuration, seven frontend, five tool-calculation, 11 static-export and one admin-header test, plus 281 retained automation checks (71, 46, 38 and 126 end-to-end). Scoped Ruff lint, Mypy across six native Python files and the Bandit gate passed, and Gitleaks found no leaks in the staged documentation diff or the 46-commit history. The retained AskMe content-search test fails under plain Node because its JSON import and generated index are missing; that failure is logged in the defect log and is not counted as a pass. Two integration suites that need a running stack were not run.

25. Retained automation and optional publication tooling

The repository retains customer-schema validation, deterministic planning/dry-run behavior, structural Elementor parsing, content/media orchestration, a centralized WordPress client, an authenticated legacy REST bridge, static export and bounded AskMe retrieval tooling. These are separate from the current native reader path and must be revalidated before reuse against another customer or environment.

The intended automation sequence is configuration → validation → reviewable plan → authenticated narrow mutation → readback → publication artifact → verification. Structural matching should fail on ambiguity rather than selecting an arbitrary widget. A successful HTTP write is not enough; read back the decisive fields and preserve unrelated document data.

The optional static exporter validates origin/path scope, rewrites nested asset references and excludes unsupported dynamic submissions. Its public content index must contain only approved public material and be released with the matching snapshot. The retained AskMe worker is bounded retrieval, not unrestricted clinical reasoning, and is not called by the current browser library.

Application authentication required by legacy bridge tooling conflicts with the native runtime's deliberate application-password disablement. Do not enable old workflows merely because their files exist. Review the intended authority, credentials, exposed routes and tests as a separate activation.

The 24 September documentation review found that a legacy static snapshot deployment still exposes the retained AskMe route publicly, independently of the native site. The native reader path still does not call it. Decommissioning that snapshot, or restricting the route's origins and adding rate limiting, is an open owner decision tracked in the roadmap's quality track and the threat model. Until it is resolved, the route is treated as a current exposure, not as a retired component.

26. Developer handoff and configuration audit

Start with wordpress-native/README.md, ACCEPTANCE.md, the typed environment loader, publication/delivery policies and the native builder. Trace one page from maintained source into Elementor data, through the scoped rendering plugin, gateway response and browser behavior. Then exercise an editor save/restore and inspect metadata/cache behavior.

Use local-first drift checks for every production change. Keep credentials out of code, test fixtures, logs and public documentation. Treat source, maintained content and editor-owned database state as separate owners; never reimport a frozen layout over newer editor changes. Rebuild Elementor CSS and invalidate the relevant cache after layout imports.

Configuration regression tests reject obvious credential material and provider identifiers in the native runtime and constrain fixed URL identifiers. Environment examples contain safe values only. The audit and scans establish the reviewed scope; retained older modules are not silently represented as having received a complete new architecture refactor.

27. Client ownership and acceptance

The client owns legal identity, approved copy, media rights, editorial verification, intended clinical/commercial use and the decision to activate optional collection or monetization. Developers own faithful implementation, permission boundaries, test evidence, recovery and truthful reporting of limitations.

The owner should compare the preserved original design with the public native site, inspect the brain hero and slower scroll, review desktop/mobile spacing, and edit representative content in WordPress. Check header/footer templates, ACF metadata, library search, saved items and all six tools. Do not change real content merely to prove editing without a recovery plan.

Engineering checks passed for the recorded release. Owner aesthetic approval remains separate. Search ranking, complete legal compliance, clinical validity, physical-device coverage, sustained concurrency and long-term uptime have not been certified.

28. Operating roadmap and change control

Near-term work for a real client is operator identification, current jurisdiction review, editorial verification, accessibility testing with intended users/devices, independent monitoring and recurring restore drills. Optional features remain off until their data, authorization, legal and operational requirements are implemented and tested.

Scale work proceeds from measured workload and delivery strategy to redundant stateful architecture and representative load/failure qualification. Additional infrastructure or paid services require explicit cost approval. A change to the workload or business model reopens relevant security, privacy, capacity and publication assumptions.

Keep release notes and the public overview synchronized with actual evidence. Do not convert a planned feature into a verified claim because code or a diagram mentions it. Preserve tested rollback assets and update recovery records as each meaningful stage completes.

The published roadmap formalizes this sequence. Phase 0 (foundation) is complete. Phase 1 establishes observability and the 99% SLO window, Phase 2 edge delivery for 10k–100k concurrent readers, Phase 3 a redundant origin and database high availability, and Phase 4 1M concurrent-reader qualification. Track L activates legal and commercial features one at a time. Track Q collects quality and CI work, including enabling the guarded workflows, pinning actions by SHA, moving CI off end-of-life Node 20, fixing the AskMe test harness, removing machine-specific paths, binding local Docker ports to loopback, an accessibility audit and resuming Dependabot. A phase is complete only when its exit criteria have evidence.

29. Reference material and maintenance entry points

The repository's native verification record, US/EU research report and capacity/availability roadmap provide the detailed supporting evidence. Configuration files contain the executable assumptions; test results identify what was actually exercised. The public overview is a scrubbed derivative of the private project record and excludes recovery secrets and infrastructure exposure.

Useful primary technical references include WordPress hardening, WordPress performance, NGINX request limiting, Google structured-data guidance and the inline regulator sources above. Research and release checks are time-bound; refresh them before a materially different client launch.

The repository documentation index links every current document, including the architecture, operations runbook and development guide. Project history is kept in the changelog and defect log; contribution and support routes are in CONTRIBUTING.md and SUPPORT.md.