Mosaic Labs Docs

Architecture, URL map, editorial state, Google Images and payment / release controls.



WEBSITE • PRODUCTS • OPERATING BOUNDARIES

This reference describes confirmed contents of the archived V12 static website and clearly separates that source state from the offline updated candidate. It does not certify production security, payment operations or release readiness.

EDITION STATUS / OWNER REVIEW COPY — NOT FOR SALE OR DEPLOYMENT

This controlled document is a factual and editorial review candidate. It must not be treated as an approved sale, investment offer, tested checkout, published policy, or evidence of features that are only planned.

1. System context and version boundaries

The original reference website is a static multipage V12 bundle distributed as a 77-entry ZIP. It contains HTML, CSS, JavaScript, JSON product metadata, icons and visual assets, PDF drafts, sitemap.xml, sitemap-images.xml, robots.txt, _redirects and _headers. This source bundle is kept unchanged; edits are performed on a separate candidate copy with a file-level diff.

The updated candidate has the same legacy pages and tools, plus a new monthly service route, reconciled digital-product descriptions, owner-selected title adjustments, 14 official indexable compositions and three controlled PDFs. It is not deployed. Static bundle contents are not evidence of identity with every resource currently served by Cloudflare or any other origin.

Related surfaces: the primary website at mosaiclabsprotocol.com and the independently hosted app.mosaiclabsprotocol.com. Data flow and deployment controls for the separate dApp were not provided and are not invented here.

2. Public divisions and offer types

Electrical: six browser-accessible free diagnostics; Google Visibility Starter Kit (digital); Google Visibility & Local SEO Sprint (digital, under the existing historical /services/ URL); and Search Growth Management (one proposed monthly managed service, three tiers). The digital products are independent self-guided workbooks/manuals rather than professional client engagements.

Crypto: Mosaic Labs Digital Asset Research — Selection & Portfolio Architecture, a historical education and research-methodology package, rather than a live ETF, custody product, token, signal or investment advisory system.

dApp: public materials describe independent read-only portfolio visibility. Do not infer that holdings are necessarily complete, wallet connectivity is universally supported, trading is possible, or a token is issued.

Docs / Whitepaper / Offer: versioned informational documents linked from existing canonical public pages. Their status must be visible and must not be confused with an approved buyer package.

3. Public URL directory

Homepage
https://mosaiclabsprotocol.com/

Electrical division
https://mosaiclabsprotocol.com/electrical/

Crypto division
https://mosaiclabsprotocol.com/crypto/

Starter Kit product page
https://mosaiclabsprotocol.com/electrical/digital-products/google-visibility-starter-kit/

Digital SEO Sprint product page (legacy /services/ route)
https://mosaiclabsprotocol.com/electrical/services/google-visibility-local-seo-sprint/

Digital Asset Research product page
https://mosaiclabsprotocol.com/crypto/research/digital-asset-research/

Monthly SEO service candidate (not live yet)
https://mosaiclabsprotocol.com/electrical/services/search-growth-management/

Free Tools hub
https://mosaiclabsprotocol.com/electrical/free-tools/

Complete Offer page
https://mosaiclabsprotocol.com/offer/

Whitepaper page
https://mosaiclabsprotocol.com/crypto/whitepaper/

Docs page
https://mosaiclabsprotocol.com/crypto/docs/

dApp
https://app.mosaiclabsprotocol.com/

Contact
https://mosaiclabsprotocol.com/contact/

4. Six Free Tools: inventory and boundaries

These six V12 routes contain actual HTML and script assets that must be preserved unchanged in this mandate. The audit of supplied static files confirms that each route has an index.html, but no authenticated live-user field test, real calculation end-to-end test or native migration acceptance has been recorded. A tool must not be described as a Google-provided rating or direct, automatic GBP/GSC integration without proof.

01 — 27-Point Google Lead Leak Audit

A staged diagnostic to capture website visibility, contact and operational observations. A native-staging subroute is explicitly marked noindex and is not the approved public migration.

Open live tool →

02 — Lead Response Revenue Calculator

Hypothetical follow-up and revenue sensitivity calculations from user-entered values, not observed business outcomes.

Open live tool →

03 — GBP Local Visibility Score

A structured Google Business Profile checklist/scoring aid, not an official Google rank score.

Open live tool →

04 — Website Conversion Score

Observations on website contact elements, mobile friction and enquiry paths.

Open live tool →

05 — Service Page SEO Score

A locally relevant service-page diagnostic for headings, intent, content and customer clarity.

Open live tool →

06 — Review Velocity Calculator

Review pace and user-entered volume analysis; no rankings, reviews or reputation outcomes guaranteed.

Open live tool →

5. Digital product architectures and release states

Starter Kit ($79 one-time planned): the R18 owner-review archive has a 245-page ebook, a 13-tab candidate workbook for 30 evidence-based checks, original templates, 20 editable field forms and an action plan. Desktop Excel and novice tests remain unexecuted; 14 legal/scope/owner questions remain open. It is not a certified commercial release.

Local SEO Sprint ($149 one-time planned): the 16 October owner-review archive has a 227-page ebook, a 16-tab Sprint Control Center, ten editable DOCX templates, a four-week plan, worked case and safety materials. It is a self-directed product. The archive states that native Excel tests were not run and the license and owner approval are pending.

Digital Asset Research ($69 one-time planned): 271-page ebook, 18-sheet workbook, historical archive and an authentic 28-page review preview. The English archive is a derived working 371-page document, not a certified translation of the original 346-page Romanian source. The historic 125 = 44+32+49 split and 60/30/10 allocation are examples from a dated research record. Investment, source-rights and legal release limitations remain.

The public V12 JSON catalogue originally identifies Sprint as type service. The controlled candidate changes only that classification and relevant language to digital_product and adds a separate disabled monthly_service definition. Checkout remains disabled for all offers.

6. Monthly service model and owner approvals

Search Growth Management is a service family with Essential $399/month, Growth $899/month and Advanced $1,999/month. One client engagement is bounded to one U.S. electrical business, one principal website/domain, one Google Business Profile and one primary local market. Client onboarding verifies authorization and access; the owner retains commercial responsibility.

The plan copy must explain month-one onboarding, monthly work ceilings, what qualifies as an existing page update or limited technical fix, review cadence, customer dependencies, communication and exclusion of ad spend, complex custom development or results guarantees. The source dossier labels the same plan prices Foundation/Momentum/Expansion under Search Presence Management. Brand and contract language cannot be treated as harmonized until an owner decision is recorded.

Three selection controls are shown on the candidate service page. They are intentionally disabled; there are no Stripe Price IDs or recurring Checkout sessions. The buttons cannot silently enroll a visitor.

7. Payment, fulfillment and recurring billing — design only

The proposed architecture begins with **Stripe Checkout**. Upon a signed server-side payment-completion event, a minimal order ledger associates session ID, SKU, version and customer email; event processing must be idempotent. Protected private files are delivered through temporary scoped links and a transactional email from Resend. A success page alone is not acceptable evidence of an authorized paid order.

For subscriptions, a distinct recurring Price ID per tier must be created after approved terms, billing interval and account requirements. Server-side records must distinguish created, paid, incomplete, past_due, canceled, renewed and failed-payment states according to the selected integration. Cancellation, refunds, taxes, failed payment, support and recurring renewal disclosure must be verified before enrollment.

The supplied ZIP contains neither accepted server-side webhook code nor approved live secrets, private fulfillment store, fulfillment tests, validated Stripe subscriptions or deployed Resend delivery. No payment systems were accessed or changed in creating this candidate. Do not publish checkout as operational or claim automated delivery.

Recommended test sequence (not executed)

- Create test-mode products and approved prices; verify one-time payments and three distinct recurring plans.
- Verify webhook signatures, duplicate events, async payments, mismatched price IDs, failed billing, renewal and cancellation.
- Verify private file access controls, expiring link behavior, resend/idempotency, unsubscribe differences and minimal customer-data retention.
- Complete owner, rights, legal, tax, refund/cancellation and support acceptance before production enablement.

8. Image indexing and search surfaces

The original V12 archive has 12 official JPEG image assets in /assets/google-images/ with homepage-linked thumbnails and an image sitemap. The approved chat supplied two additional distinct visual materials: the actual three-tier SEO service image and a public website presentation screenshot. The candidate preserves the 12 V12 JPEGs byte-for-byte and adds two PNGs for a total of 14 distinct official indexed compositions. For traceability, original owner uploads matching the original 12 compositions are also packed separately without altering their V12 presentation counterparts.

The original assets/research.webp remains an HTML display asset but has a targeted proposed X-Robots-Tag: noindex to address the previously requested image removal. Other pre-existing display-only .webp assets receive the same index-exclusion intent. Official imagery in /assets/google-images/ must remain accessible and eligible to be crawled.

sitemap.xml preserves the existing canonical routes and adds only the service route. sitemap-images.xml associates all 14 official image URLs with the homepage. This expresses publisher intent; Google decides actual indexing and landing-page presentation. After authorized deployment, verify HTTP 200 image responses, correct content types, canonical page, no blocking robots rules, Cloudflare header behavior and Search Console observations.

Image migration and legacy Google results

Temporary Search Console removals are not permanent removals. When an old image must never return to results, it must become non-indexable permanently: remove it and serve 404/410, or return a crawlable X-Robots-Tag: noindex. Do not assume deleting a sitemap or a Cloudflare Pages project removes image results from Google.

The owner supplied an explicit request to retire the Google Images URL /assets/research.webp, while retaining the research page itself. The candidate’s targeted image response header is intended to preserve visual rendering but exclude that historic URL. It cannot be certified until the host is deployed and headers are checked.

9. Static routing, SEO and headers

V12’s _redirects file must be preserved for historic routes; redirect chains and expected HTTP response codes should be validated only after authorized deployment. _headers carries content and framing policies, a native-staging noindex rule and draft PDF noindex. The candidate appends image-specific exclusion headers without changing global security headers. No other header policy or tool behavior is altered.

The primary sitemap points to indexable canonical pages and the image sitemap to owner-approved images. robots.txt remains the historic source file. These documents aid discovery but do not guarantee Google inclusion. Deleting a Search Console property or sitemap entry does not deindex existing pages; Google Search Console reports may continue to show redirects, valid canonicals and retired 404s without indicating a site malfunction.

The homepage has one scope-limited CSS text-align change on the short cyan eyebrow directly under MOSAIC LABS PROTOCOL. Global layout, brand mark, navigation and responsive rules are preserved. Layout and interaction regressions remain an acceptance gate; a static file audit is not a substitute for real mobile/desktop browsing.

10. Candidate document URLs and source preservation

Draft source PDFs remain present in the archive. New controlled documents are placed separately under /docs/ with stable review filenames. The public documentation pages in the candidate link to the new versions. The previous single-page Offer source document remains in the archive but the Offer page now links to the more complete multi-page Offer review PDF.

Candidate document paths (not live until deployment): /docs/MOSAIC_WHITEPAPER_REVIEW.pdf, /docs/MOSAIC_DOCS_REVIEW.pdf and /docs/MOSAIC_COMPLETE_OFFER_REVIEW.pdf.

11. Verification, security and release gates

A local archive verification can prove that assets exist, image references resolve, XML parses, URLs are internally consistent, canonical links are declared, source files remain byte-identical where not authorized, PDF annotations contain hyperlinks, and new checkout controls are disabled. It cannot prove runtime behavior behind Cloudflare, external service availability in every region, network security, payment processing or accessibility with assistive technology.

Release gate order: (1) final owner approval of naming, terms and content; (2) solve product rights, finance/legal and application tests; (3) complete desktop/mobile/browser checks on a true staging host; (4) test the host’s redirects, images, security headers and permanent removals; (5) configure and validate Stripe–Resend if separately authorized; (6) review the diff and approve deployment; (7) submit sitemaps and observe Google recrawling. No step is deemed complete solely by being listed here.

12. Source trace and document limitations

Evidence used: 79-page Master Execution Dossier V2 dated 14 October 2026; original V12 ZIP (77 entries); Starter Kit R18 owner-review handoff; SEO Sprint release-status memo dated 16 October 2026; Mosaic Labs Digital Asset Research release-hold and source notes; local static files, PDFs and online read-only URL checks.

No independent security penetration test, native desktop Excel use, publication permission for the source M30 archive, verified Stripe transaction, Resend production email, dApp backend audit or complete factual certification is claimed. These docs deliberately retain the correct negative statements until acceptance evidence exists.