SECURITY ASSESSMENT REPORT

Éclair CMS

Two source-code security audits (1.10.2 and 1.10.3) and the security releases 1.10.3 and 1.10.4

DEFENSIVE REPORT · NO EXPLOITATION DETAIL

PROJECT
Éclair CMS — open-source PHP CMS

REPOSITORY
github.com/Jacckisback/eclair-cms

AUDIT 1
Éclair 1.10.2 · commit 796bc8d · complete

AUDIT 2
Éclair 1.10.3 · commit 001d224 · deep profile, partial pass

PATCHED RELEASES
1.10.3 (2026-11-03) · 1.10.4 (2026-11-04)

METHODOLOGY
Cloudflare security-audit skill

EXECUTION
Source reading and local tests in an OS-enforced sandbox

REPORT DATE
4 November 2026

This report describes, for each issue, what was wrong and how it was fixed. It deliberately contains no reproduction steps, payloads, request examples or exploitation instructions. Technical identifiers (file paths, function names) are given as recorded in the audit records.

1 Executive summary

Éclair CMS was subjected to two successive source-code security audits using the Cloudflare *security-audit* skill, each followed by a dedicated security release. All analysis was performed by reading the source and running local tests inside an OS-enforced sandbox, on dummy data, without network access and without probing any deployed site.

148
agents across both audits and the fix work

21
confirmed flaws in total (15 + 6), all fixed

0
critical findings in either audit

964
automated checks in 1.10.4 (source and signed archive)

Stage	Target	Outcome
Audit 1	1.10.2 · 796bc8d · standard profile · complete	86 agents, 50 coverage units, 24 records: 15 confirmed (3 medium, 12 low, 0 high, 0 critical), 8 needs-validation, 1 rejected.
Release 1.10.3	Security release	All 15 flaws and all 8 leads fixed in 5 fix lots; more than 50 hardening items; automated checks raised from 620 to 918; real 1.10.2 → 1.10.3 update tested from GitHub.
Audit 2	1.10.3 · 001d224 · deep profile · partial pass	54 agents, 71 coverage units (50 covered, 15 candidate, 6 deferred), 13 distinct candidates: 6 confirmed (1 high, 1 medium, 4 low), 6 needs-validation, 1 unresolved. The 1.10.3 fixes held.
Release 1.10.4	Security release	13 fixes, each with a regression test; 964 automated checks on source and signed archive; Ed25519-signed announcement (legacy and v2 signatures verified); archive 621 KB.

Key messages

- **No critical finding in either audit.** Audit 1 found no high-severity issue in 1.10.2. Audit 2 found a single high-severity issue, located in the maintainer's release tooling (`tools/build.php`) and never in an installed site.
- **Every confirmed flaw and every lead has been closed in code.** The 15 flaws and 8 leads of Audit 1 were fixed in 1.10.3; the 6 flaws, 6 host-dependent leads and the unresolved candidate of Audit 2 were addressed in 1.10.4.
- **Fixes were independently re-examined.** Audit 2 re-planned every earlier record against the changed source with fresh agents; the Audit 1 flaws came back as fixed.
- **The core defensive design held in both audits:** signed updates, Argon2id password hashing, strict sessions, CSRF with origin checks, allowlist content sanitisation, a strict Content Security Policy, OAuth with PKCE S256 for the MCP connector and an authenticated-encryption vault.
- **Audit 2 is a partial pass.** It was shortened at the owner's request to limit cost: 6 coverage units were deferred and no final coverage critic ran. No claim of complete coverage is made for 1.10.3 or 1.10.4.

> **Positioning.** Éclair ships with no plugins, no themes from third parties and no third-party dependencies. This removes, by construction, the extension ecosystem that accounts for most published vulnerabilities in comparable CMS platforms (see section 8). It does not by itself make the software free of defects, as the findings in this report show.

2 Methodology and sandbox

Method

Both audits followed the Cloudflare *security-audit* skill (github.com/cloudflare/security-audit-skill). The skill organises a full audit into reconnaissance, coverage planning into units (surface × trust boundary × subsystem × attack class), waves of independent "hunter" agents, coverage critics, and independent verification of every candidate by fresh agents before anything is reported. Each candidate ends with one of three dispositions: *confirmed*, *needs validation* (a real code trace whose impact depends on a fact that cannot be observed offline) or *rejected*.

- **Audit 1** used the *standard* profile, with separate Phase 3 (candidate verification) and Phase 5 (final record verification) verifiers.
- **Audit 2** used the *deep* profile for hunting, which adds a *lifecycle mode* dimension to each coverage unit (update, restore, rollback, reinstall). At the owner's request, validation then merged Phases 3 and 5 under the quick-profile rule: one fresh, independent verifier per candidate, which also performed the final record checks.

Scope

For both audits: `eclair/`, `tools/`, `dist/latest.json`, `.github/workflows/` and `tests/` (local fixtures and harnesses only), on a clean working tree of the audited commit. Agents worked on an archive copy of the audited commit, never on the developer's working copy.

OS-enforced sandbox

Every agent invocation ran in a sandbox enforced by the operating system rather than by convention:

Namespaces	Separate mount, network, PID and IPC namespaces; no external network, isolated loopback only.
System-call filter	seccomp policy forbidding, among others, namespace creation, mounts, `ptrace`, `bpf` and `AF_VSOCK`.
Privileges	`setpriv` to user `nobody`, no capabilities; reduced allowlisted environment.
Resource limits	cgroups: 1 GiB memory, 1 CPU, 64 tasks.
File system	Root and target read-only; writes limited to each agent's own scratch directory.
Data	Dummy data and dummy keys only; no real credentials, no real user data.
Deployed systems	No probing of any deployed site and no call to any external service.

Agent effort

Activity	Audit 1 (86 agents)	Audit 2 (54 agents)
Reconnaissance	4	5

Activity	Audit 1 (86 agents)	Audit 2 (54 agents)
Hunters	22, in 5 waves	34 (33 + 1 replacement), in 1 wave
Coverage critics	5 post-wave critics + 1 final clean critic	1 post-wave critic, stopped before any result at the owner's cost request
Verifiers	26 Phase 3 + 28 Phase 5	14 (merged Phases 3/5), including 1 re-run

Together with the fix work for the two releases, 148 agents were used in total.

3 Audit 1 — Éclair 1.10.2

Target	Éclair 1.10.2, commit `796bc8d`
Profile	Standard, full audit mode; no cap set by the owner
Coverage	50 units; 5 post-wave critics, then a separate final critic that came back clean (no missing unit, no reassignment); 0 blocked, 0 deferred, 0 out of scope
Previous runs	None in the skill's format
Status	COMPLETE

24
final records

15
confirmed: 3 medium, 12 low

8
needs validation (no severity)

1
rejected

The audit concluded that the code is defensive overall: cryptography, update signing, content sanitisation, sessions, CSRF and OAuth hold on most of their paths. The most significant issues were sensitive actions becoming available again without a password, editor-controlled redirects on domain-validation addresses, and an uncapped store that could block editing. Most of the others concern race conditions, derived copies (cache, index, logs) and quotas that did not keep their documented promise.

Confirmed findings (15)

#	Severity	What was wrong	Fix in 1.10.3
1.1	MEDIUM	"Log out my other devices" and "Log out everyone" re-ran the login, renewing the 15-minute password confirmation and the 48-hour session lifetime without a password.	These actions now re-key the session without renewing the password confirmation or the session lifetime.
1.2	MEDIUM	Editors could create stored redirects on operator-reserved addresses (`/.well-known/` , including domain-validation paths, and `/_eclair/`).	These namespaces can no longer be redirected by editors; existing redirects there are disabled. Admins can still redirect other `/.well-known/` addresses such as `change-password` .
1.3	MEDIUM	The content index and search store had no aggregate cap; mass-created drafts could make saving, deleting and visitor search fail.	Store capped at 3,000 contents; the AI stops creating drafts at 500; the search index keeps only published text, up to 16 MB.
1.4	LOW	Deleting an author account left their public profile in the static page cache (author page, bylines, feeds, exports).	Account deletion now clears the page cache.

#	Severity	What was wrong	Fix in 1.10.3
1.5	LOW	The AI assistant or an MCP client could change live content without a proposal when the item was published or unpublished during its turn.	Direct edits apply only to an item that is still a draft at write time; otherwise a proposal is created.
1.6	LOW	The trusted-device cookie survived password changes, forced logouts and resets, keeping its own login-attempt allowance.	Device trust is now tied to the password and the account's session version; every revocation makes other browsers forget it.
1.7	LOW	Anonymous logins created one persistent limiter file per submitted e-mail, and the Security page loaded them all into memory.	Account limiters use a fixed number of buckets; the Security page can no longer be made to fail by them.
1.8	LOW	When `ECLAIR_DATA` ended with a slash, the backup skip list failed open and exported the vault together with its key.	Archive paths are derived independently of the spelling; backups never contain the vault, sessions, limits or update state.
1.9	LOW	A page-cache write completing after a purge could keep retracted content public (render/purge race).	The page cache is versioned by purge generation; a page rendered during an unpublish can no longer stay cached.
1.10	LOW	An unlocked index rebuild (translation link or admin reindex) could republish an unpublished draft, because public renderers trusted the index status.	Rebuilds are serialised with writes, and public outputs check each item's own saved status.
1.11	LOW	Anonymous login failures could push real security events out of the security log.	Anonymous failures are kept in their own file; logs roll over to a previous segment instead of losing their older half.
1.12	LOW	The contact form accepted an address that could add a header line to the owner's notification e-mail.	Addresses with quotes, backslashes, spaces or line breaks are refused everywhere.
1.13	LOW	A rate-limit block restarted the accounting window, so contact and MCP write quotas refilled early (about 4× their ceiling) and login limits accepted a second batch.	Rate limits no longer refill before the window in which they were exhausted has ended.
1.14	LOW	A concurrent request of the same session could re-create a session file destroyed by logout.	Session destruction and write-back now run under a lock; logout can no longer be undone.
1.15	LOW	Signed update announcements carried no serial or expiry, so whoever controlled the channel file could replay an older one to hide updates or re-offer a withdrawn release.	Announcements carry a signed serial and expiry (v2 signature); older announcements are refused and reported.

Leads requiring validation (8)

These have a real code trace but depend on a deployment or browser fact that cannot be observed offline. They carry no severity. All eight were nevertheless closed in code in 1.10.3.

#	Lead	Closed in 1.10.3
L1	Uncapped editor-written stores (redirects, media, categories) processed in full by anonymous routes, with each redirect hit rewriting the whole redirect table under the global data lock.	Redirect hits no longer rewrite the list under the site-wide lock; redirects (5,000), images (10,000) and categories (1,000) are capped.
L2	The two-at-a-time cap on password checks was lifted when a login lock file could not be opened.	An unopenable lock file no longer lifts the limit.
L3	Trusted-device and OAuth logins shared an unprioritised two-slot verification pool, so a large failure flood could delay the owner's login.	One check is reserved for usual devices and signed-in users.
L4	Markdown exports (`.md`, `llms.txt`, `llms-full.txt`) did not fully neutralise markup and links typed as text.	Exports can no longer carry HTML, images or links typed as text; real links keep their checked address.
L5	The Google Drive connection and backup binding were not tied to the admin session that passed the step-up, and outlived the connecting admin.	The connection is tied to the starting admin session and removed if that admin is deleted or demoted.
L6	MCP OAuth consent for a self-registered client was a single click, without protection against window substitution (double-clickjacking).	Apps other than Claude or ChatGPT always require the password; the "Autoriser" button is protected against double-clickjacking.
L7	The data store's critical sections ran unlocked when `data/.lock` could not be opened.	The lock is never bypassed: Éclair recreates it or refuses to save with a clear message.
L8	Unsaved editor drafts in browser storage outlived logout and could be offered to the next account in the same browser profile.	Backup copies belong to the signed-in account and are erased at logout.

Rejected (1): a proposed issue on re-enrolling the TOTP second factor without re-authentication was refuted by verification. 1.10.3 nonetheless tightened 2FA re-enrolment (disable with the current password, then enable again).

4 Release 1.10.3 — security patches (3 November 2026)

15 + 8 flaws and leads fixed	**5** fix lots (A, B1, B2, C1, C2)	**50+** hardening recommendations applied	**620 → 918** automated checks

Security release fixing every issue found by the full audit of 1.10.2. A real update from 1.10.2 to 1.10.3 was tested from GitHub. Sites on 1.10.0–1.10.2 update in one click from Administration › Updates.

Security fixes — medium

- Device and global logout no longer renew the password confirmation or session lifetime.
- `/.well-known/` and `/_eclair/` can no longer be redirected by editors.
- Content store capped (3,000 contents, AI drafts stop at 500, search index limited to published text up to 16 MB).

Security fixes — low

- Account deletion clears the page cache.
- AI/MCP edit directly only items that are still drafts.
- Usual-device cookies tied to the password and session version.
- No per-e-mail limiter files; the Security page stays available.
- Backups never include vault, sessions, limits or update state.
- Render/purge race closed in the page cache.
- Index rebuilds cannot republish drafts; public outputs check each item's status.
- Anonymous failures can no longer evict security-log events.
- Contact-form header injection closed.
- Rate limits no longer refill early.
- Logout can no longer be undone by a concurrent request.
- Signed serial and expiry on update announcements.

Leads closed in code

- Login lock fallback no longer lifts the concurrency limit; a check is reserved for usual devices.
- The data lock is never bypassed.
- Redirect hits no longer rewrite the list under the global lock; redirects, images and categories capped.
- Markdown exports neutralise markup typed as text.
- Google Drive connection tied to the starting admin session.
- MCP consent: password for unknown apps; double-clickjacking protection.
- Editor backup copies bound to the account and erased at logout.

Notable changes

- 2FA cannot be re-enrolled while enabled; setup codes limited to 10 per 15 minutes.
- FTP reset refuses the example password from the guides and future-dated reset files.
- Updates page shows "latest version" only after a fresh signed announcement; a warning appears after 180 days without renewal.
- Security page shows real events and anonymous failures separately.
- `media/` no longer serves `.js` or `.mjs`.
- On storage without working file locks, saving is refused with a clear message.
- Installer re-run keeps existing keys and never deletes existing accounts or content.
- For maintainers: dual signature on `dist/latest.json` (legacy + v2 with serial and 180-day expiry); signed builds only from a Git checkout; new regression suites `tests/sec-*.php`.

5 Audit 2 — Éclair 1.10.3

Partial pass. This run was launched with the deep profile and then shortened at the owner's explicit request to limit token cost. Coverage stopped after the first hunting wave, the post-wave critic was stopped before producing any result, no final coverage critic ran, and validation merged Phases 3 and 5 under the quick-profile rule. Six accepted wave-2 units were deferred and never hunted. **No claim of complete coverage is made.**

Target	Éclair 1.10.3, commit `001d224`
Profile	Deep (hunting); validation folded to quick rule (Phases 3 and 5 merged)
Agents	54 of a planning ceiling of 150; the gap reflects the owner-requested early stop, not budget exhaustion
Previous run	Audit 1 on 1.10.2. Nothing was carried over unchanged: every earlier unit and record was compared with the current source and re-examined by a fresh agent
Status	PARTIAL PASS

Lifecycle modes exercised

The deep profile adds a lifecycle dimension to each coverage unit. Modes exercised:

Normal operation
A running 1.10.3 site.

Post-update
First requests after updating from 1.10.0–1.10.2, with legacy data states.

Restore
Backup restored over the site via FTP.

Install / update / rollback
Install, reinstall, one-click update and rollbacks.

Several Audit 2 findings arise precisely from these lifecycle modes (legacy data after update, restore, rollback, reinstall).

71
coverage units: 50 covered, 15 candidate, 6 deferred

13
distinct candidates

6
confirmed: 1 high, 1 medium, 4 low

266
hardening notes from hunters

The 1.10.3 fixes held. 1.10.3 changed the source of each of the 15 confirmed flaws and 8 leads of Audit 1. All of these were re-planned against the changed source (14 former confirmed flaws, 6 former leads, the former rejected lead and 29 previously covered units) and re-examined by fresh agents. In this run, the Audit 1 flaws came back as fixed. The records below are new issues specific to 1.10.3.

Results (13 candidates)

#	Disposition	What was wrong	Fix in 1.10.4
2.1	HIGH	**Release tooling only (`tools/build.php`), never an installed site.** The withdrawal command `--resign=<version>` signed whatever archive was in `dist/`, bound only by its internal version string, so an archive substituted by someone without the signing key could have been re-announced under the maintainer's signature.	The command now signs only an archive whose SHA-256 and size match a key-verified announcement in the Git history of `dist/latest.json`, or a SHA-256 the maintainer supplies explicitly.
2.2	MEDIUM	Google Drive connections made before 1.10.3 had no recorded binder and escaped the new guard, so a demoted or deleted admin could keep receiving full site backups.	Such connections are revoked; the Security page asks an administrator to reconnect.
2.3	LOW	An index rebuild could republish content deleted under 1.10.0–1.10.2 when a failed delete had left its file behind.	A file the index does not list comes back as a draft for review.
2.4	LOW	Restoring a backup over the site could revive sessions and MCP connections cut by "Log out everyone".	Revoked sessions and MCP connections are now deleted, not just refused, on every revocation event.
2.5	LOW	A failed write after an unpublish or delete skipped the cache purge, so the page cache kept serving retracted content.	The cache is purged on every path once writing has started, even when a later write fails.
2.6	LOW	Manual rollback unconditionally overwrote `.htaccess` and directory-protection files changed after the update.	Changed files are kept, their previous version is saved in `cache/update`, and the page says so.
2.7	NEEDS VALIDATION	Image generations rejected by provider moderation, or failing after submission, escaped the monthly AI budget (depends on third-party billing).	They now count toward the monthly AI budget.
2.8	NEEDS VALIDATION	After loss of `data/site.php` alone, the installer could reopen over existing accounts, keys and vault (depends on hosting file events).	The installer stays locked until `data/install-window.php` is deleted via FTP.
2.9	NEEDS VALIDATION	Anonymous MCP registration and token calls created one persistent rate-limit file per caller network, with no site-wide bound (depends on host disk/inode quota).	Counters grouped into at most 4,096 per kind.
2.10	NEEDS VALIDATION	Anonymous share-image generation decoded the full original without the upload memory guard (depends on the host's GD build and memory limit).	Decoding beyond the memory limit is avoided; the page renders and is cached with the original as share image.

#	Disposition	What was wrong	Fix in 1.10.4
2.11	NEEDS VALIDATION	Sensitive data files passed through a temporary file with the default umask mode before being restricted to `0600` (depends on host account isolation).	Data files and backup archives are created private from the start; the data folder is tightened to `0700` where allowed.
2.12	NEEDS VALIDATION	Pre-1.10.3 editor backup drafts from another install on the same origin, or a removed account, could be adopted by a single-account site (depends on browser and deployment).	Unattributed copies of never-saved content are erased; copies of existing content are kept.
2.13	UNRESOLVED	Redirect chain-collapse could rewrite a target onto a reserved address. Two fresh verifiers completed their checks, but the platform's content filter withheld their final output both times, so no independent disposition exists.	Fixed as a precaution: editors can no longer change the target of an administrator's `/.well-known/` redirect through chain collapse.

Deferred coverage (6 units, never hunted)

Declared gaps, not evidence of safety: (1) per-item content footprint and revision retention; (2) backup archives at rest; (3) client identity and scheme from proxy headers; (4) cardinality of anonymous rate-limit counters; (5) third-party runtime artefacts (self-hosted font download without pinning); (6) quotas of the MCP OAuth endpoints shared across a platform's users.

6 Release 1.10.4 — security patches (4 November 2026)

13	964	Ed25519	621 KB
fixes, each with a regression test	automated checks on source and signed archive	signed announcement, legacy + v2 verified	release archive

Security release closing every result of Audit 2. Sites on 1.10.0–1.10.3 update in one click from Administration › Updates, after a backup. The update is signed and verified before anything is touched, with automatic rollback.

Class	Fix
HIGH release tooling	`tools/build.php --resign=<version>` signs only an archive already signed for that version (SHA-256 and size matched against a key-verified announcement in the Git history of `dist/latest.json`) or one whose SHA-256 the maintainer supplies in `ECLAIR_EXPECT_SHA256`.
MEDIUM	A Google Drive connection made before 1.10.3 (or redone by 1.10.2) is removed until an administrator reconnects it.
LOW	Restoring a backup no longer revives sessions or MCP connections cut by "Log out everyone", "Log out my other devices", a new password, a password reset (admin or FTP) or account deletion.
LOW	Unpublishing or deleting content clears the page cache even when a later write fails.
LOW	Rollback no longer overwrites a `.htaccess` or folder-protection file changed since the update; it is kept and its previous version saved.
LOW	Rebuilding the index can no longer republish content deleted before 1.10.3 whose file was left behind.
PRECAUTION	Editors can no longer change the target of an administrator's `/.well-known/` redirect via redirect chain collapse.
HARDENING	Data files (accounts, keys, vault) and backup archives created private from the start; data folder tightened to `0700` where the host allows.
HARDENING	If `data/site.php` is lost, the installer stays locked until `data/install-window.php` is deleted via FTP; reinstalling keeps accounts, keys and content.
HARDENING	Images rejected by the provider's safety filter (FLUX, Ideogram) or lost after submission count toward the monthly AI budget.
HARDENING	The share image of a very large photo is no longer decoded beyond the host's memory limit.
HARDENING	MCP anonymous registration and token counters grouped into at most 4,096 per kind.
HARDENING	Backup copies of never-saved content left by 1.10.2 (with no site or account) are erased instead of being given to the site's only account.

Still guaranteed. Content does not move and the public site stays byte-for-byte identical; automated checks run on the source and on the archive, with one regression test per fix; the 1.10.0–1.10.3 → 1.10.4 update is Ed25519-signed and verified before anything is touched, with automatic rollback.

7 Positive security controls verified

Beyond the findings, the audits recorded defensive patterns that were verified as sound in the code. These are the controls on which the findings above were measured.

Signed updates

Ed25519-signed announcements with a context-bound v2 signature carrying a serial and an expiry (since 1.10.3); SHA-256 and size checks; host allowlist without redirect following; manifest and PHP syntax verified at extraction; automatic rollback after a failed health check.

Password storage and login

Argon2id hashing; dummy hashing for unknown accounts; per-network and per-account limits; a bounded pool of simultaneous checks with a slot reserved for trusted devices.

TOTP with replay protection

TOTP and recovery codes are single-use, including under concurrent requests; the pending 2FA state is bound to the password and session version; FTP recovery bounded to 24 hours.

Strict, versioned sessions

Strict mode, HttpOnly, SameSite=Strict, Secure over HTTPS; a session version that revokes every device on password change; browser binding; idle and absolute expiry.

CSRF and origin checks

Session token plus origin check on every admin write; double-submit protection on the login form.

Content Security Policy

Strict CSP on the administration and on `/media` (no default source, sandboxed, `nosniff`, explicit image type); text exports served with a "none" policy.

Allowlist sanitiser

Content is sanitised by allowlist with blocks rebuilt from validated data; uploaded images re-encoded by GD.

MCP OAuth

Signed stateless `client_id`, exact redirect URI, mandatory PKCE S256, single-use hashed codes, rotating refresh tokens with reuse detection, non-frameable consent page.

Authenticated-encryption vault

Third-party API keys and tokens are kept in a vault sealed with authenticated encryption and excluded from backups.

Inert 0600 data files

Data files are PHP files guarded against direct execution output, written fail-closed under the store lock and restricted to mode `0600` (created private from the start since 1.10.4).

8 Comparison context

According to Patchstack's *State of WordPress Security in 2026*, 11,334 new vulnerabilities were published in the WordPress ecosystem in 2025, approximately 91% of them in plugins.

11,334	~91%	0
new WordPress-ecosystem vulnerabilities in 2025 (Patchstack)	of them located in plugins	plugins, third-party themes or third-party dependencies in Éclair

Éclair's architecture has no plugin or theme marketplace and no third-party package dependencies (no `composer.json` or `package.json`). The attack surface that produces the large majority of ecosystem vulnerabilities in comparable platforms is therefore absent by design, and the whole code base can be audited as one unit, as was done here.

> **Reading this comparison correctly.** The figures above describe an ecosystem, not a single product, and are not a like-for-like comparison. They explain why the absence of extensions matters; they do not rank Éclair against any other CMS. The 21 confirmed flaws in this report show that first-party code still needs regular, independent review.

9 Limitations and recommendations

Limitations

- **Audit 2 is a partial pass.** Six accepted coverage units were deferred at the owner's cost request and never hunted; the post-wave critic was stopped before producing any result and no final coverage critic ran. No complete-coverage claim is made for 1.10.3 or 1.10.4.
- **One candidate has no independent disposition.** The redirect chain-collapse candidate was fixed as a precaution in 1.10.4, but its verification output was withheld twice by the platform's content filter.
- **Host-dependent leads were not observed on real hosting.** The needs-validation items depend on hosting, browser or third-party billing facts that a sandbox cannot reproduce; they were fixed in code but not validated on a deployment.
- **Source and local testing only.** No deployed site was probed, by design; deployment configuration (web-server rules, proxy settings, file permissions) remains the operator's responsibility.
- **The 1.10.4 fixes have not yet been audited.** They are covered by regression tests, but no audit run has been performed on 1.10.4.

Recommendations

1. Resume Audit 2: hunt the six deferred units and run a clean final coverage critic, then audit 1.10.4 itself.
2. Re-run independent validation of the redirect chain-collapse candidate with a verifier whose output is not truncated.

3. Validate the host-dependent items on a controlled staging copy (file permissions and umask, GD build and memory limit, inode quota, browser storage behaviour) — never on a production site.
4. Address the main hardening themes from the 266 hunter notes: ship web-server deny rules for Nginx alongside the front-controller rewrite and probe them at install and after updates; encrypt backup archives with a key kept outside the archive; derive the vault key from a dedicated sub-key; pin self-hosted font downloads by version and SHA-256; trust forwarded-scheme headers only under the proxy setting.
5. Enable branch protection on the published repository (it lowers the release-tooling risk class of finding 2.1) and re-sign the update channel before its 180-day expiry.
6. Site operators: update to 1.10.4 after a backup, reconnect Google Drive if prompted, and use Nginx rules from the current guide.

Éclair CMS — Security Report · Audits of 1.10.2 (commit 796bc8d, complete) and 1.10.3 (commit 001d224, deep profile, partial pass) · Security releases 1.10.3 and 1.10.4 · Methodology: Cloudflare security-audit skill (github.com/cloudflare/security-audit-skill) · 148 agents in total · Source reading and local tests in an OS-enforced sandbox, dummy data only, no probing of deployed sites · Project: github.com/Jacckisback/eclair-cms · 4 November 2026.
This report is defensive by design: it describes what was wrong and how it was fixed, and contains no reproduction steps, payloads or exploitation instructions.