

EYEBALLMONITOR



Seeing availability. Preserving evidence.

Technical & strategic white paper

Product foundation, network expansion and the path to explainable incident intelligence.

ENGLISH EDITION | v1.2

04 October 2026

Public edition.

Real browsers + diverse network paths + evidence history + explainable intelligence

Reader guide & contents

For customers, partners and campaign supporters

Status convention: CURRENT describes current scope; OWNER-CONFIRMED identifies first-party deployment and test reports; CODE-REVIEWED identifies focused implementation checks; PLANNED and REFERENCE DESIGN are not deployed guarantees. English and Bulgarian editions have aligned scope. This is not an independent audit, financial prospectus, binding delivery calendar or legal opinion.

01 Executive summary	3
02 The problem and the observation gap	4
03 Why ISP and mobile visibility matters	5
04 How EyeballMonitor works	6
05 Real browsers and visual evidence	7
06 Chain of custody and evidence provenance	8
07 Managed distributed network model	9
08 Current network vs target expansion	10
09 Network map and operator targets	11
10 Incident timeline and evidence lifecycle	12
11 Customer use cases and bounded value	13
12 Controlled browser scenarios	14
13 Third-Party Dependency Intelligence	15
14 Discovery and explainable correlation	16
15 Incident Agent and autonomous evidence	17
16 Security, privacy and customer boundaries	18
17 Limitations, non-claims and validation	19
18 Roadmap and delivery gates	20
19 Commercial model and campaign context	21
20 FalcoSoft ITS and delivery accountability	22
21 Appendix: terminology and operating controls	23
References and document basis	24

01 / Executive summary

From server availability to observable user experience

EyeBallMonitor is a website-monitoring platform built around a simple question: what did a browser actually observe from a particular network path at a particular time? It combines availability checks, regional observations, browser rendering and visual evidence. The objective is to give operations teams a shared factual starting point when a site appears healthy centrally but fails for part of its audience.

The public product already presents a monitoring foundation: uptime, latency and HTTP results; browser screenshots and device profiles; incident history; customer workspaces; and Full Website Check with DAX technical analysis. These are product capabilities described by the website and supplied demonstrations, not an independent certification of every production workflow. [1, 2]

The owner reports four operating datacenter-hosted locations and 28 previously tested mobile mini-PC/SIM-router points, now suspended for lack of operating funds. The first funded mobile stage targets 80 agents in total, not 80 additions. Dependency intelligence, explainable correlation and a controlled Incident Agent remain roadmap work. The 180-country catalog is a long-term plan, not deployed coverage. [2, 3]

The project is a staged, multi-year undertaking, not a promise to complete a global network and advanced incident intelligence within a few months. The first mobile rollout is one step. Progress must follow verified access paths, reliable evidence, operational capacity and sustainable revenue.

Core thesis: an HTTP success is one observation, not a complete account of user experience.

02 / The problem and the observation gap

One website can have several different availability realities

A useful central probe can confirm that a service answered from its own location. It cannot establish that every resolver, operator, CDN edge, browser or application dependency behaved the same way. Monitoring from a datacenter remains valuable; EyeballMonitor is intended to complement it with additional perspectives rather than dismiss it.

Partial failures can arise before the HTTP request, while connecting securely, while receiving the response or after the document arrives. A page can return 200 and still show an application error, a challenge screen, an unusable layout or a checkout that cannot complete. The monitoring question must therefore state both the path and the success condition.

Failure layer	What to observe
DNS / connectivity	Resolution failure, timeout, connection error; distinguish unavailable data from a confirmed site failure.
TLS / HTTP	Certificate validation, redirect chain, status code, response time and error class.
Browser / application	Rendered content, expected elements, challenges and permitted read-only journey checks.
Dependency / region	Affected routes, third-party components and time-aligned evidence; correlation is not proof of cause.

The evidence should answer where, when and what failed before attempting to answer why.

03 / Why ISP and mobile visibility matters

Operator diversity adds information that geography alone cannot

Two users in the same city may use different access networks, resolvers, peering relationships, address families and CDN edges. A mobile connection can add carrier-grade NAT, radio variability and different routing policies. These differences matter when an incident affects only part of the audience. Geography is one dimension; the independently usable access path is another.

The architecture aims to obtain observations through actual fixed-ISP and mobile access where a connection has been provisioned and verified. Running a browser on a hosted server in a country provides a hosted perspective, not automatically a consumer or mobile perspective. The product must preserve this distinction even when both types of point are operationally useful.

Path diversity is valuable only when it is real. Multiple agents can share an upstream provider, exit address or physical access line. Country and operator targets should therefore be evaluated for independence, repeatability and relevance to customer traffic, not simply counted. A public ASN lookup is supporting metadata, not a substitute for a validated access inventory.


The intended outcome is a more precise statement: this site failed under these conditions through this verified path, while these comparison paths behaved differently. That is stronger operational evidence than an unsupported statement that a whole country or all users were down.

04 / How EyeballMonitor works

Separate orchestration, observation and interpretation

A customer adds a site, selects a monitor policy and reviews route-specific results. Lightweight uptime checks identify availability changes; browser jobs add visual evidence. Full Website Check combines browser results with DAX technical diagnostics. The public product includes customer workspaces, history and alert settings. [1, 2]

The diagram is a logical reference architecture, not a disclosure of production server topology or a claim that a particular queue, database or storage product is installed. A control plane manages customer scope, policies and jobs. Agents execute bounded checks and return structured observations. Evidence storage preserves artifacts, while incident logic and delivery workers turn results into customer-visible events.



Reference design requirement: a failed notification worker must not erase an incident or block evidence ingestion.

05 / Real browsers and visual evidence

Make the conditions of every observation explicit

Each check should preserve the target, route identity, execution time, software version, device profile, timeout and outcome. A redirect should retain both requested and final URLs. Latency needs a named definition: request duration, navigation duration and render readiness are not interchangeable. Timeout values must not be reported as measured successful response times.

Desktop, Android, iPhone and tablet views provide useful layout perspectives. Browser emulation can set viewport, device scale, user agent and touch-related behavior. It does not turn the host into a physical handset or verify the handset operating system, modem, battery, mobile radio or a specific native browser build. [8]

The success policy should be explicit: HTTP reachability, visible expected content and an authorized journey are distinct tests. Challenges, cookie dialogs and consent screens should be recorded rather than silently bypassed. A browser navigation failure may prevent a normal page screenshot; the failure remains useful evidence when its exact error and context are retained.

Comparisons should account for time skew, different resolvers, caches, content personalization and agent load. A route may be unreachable while the monitoring agent itself is healthy; an unhealthy agent may be unable to say anything reliable about the site. Missing data should stay unknown.

06 / Chain of custody and evidence provenance

A screenshot becomes useful when its context is preserved

A screenshot records a visual observation, not universal availability or proven responsibility. Its record preserves time, location, requested/final URL, HTTP outcome, device profile and check context. The owner confirms this metadata is collected. Failed navigation or capture must remain distinguishable from a successfully rendered page.

A focused code inspection found SHA-256 calculation of image bytes in the Poland, Milan and London agents and hash-linked evidence-journal entries in the hub. The upload-initialization API checks the submitted hash format. A complete persisted-file verification flow was not established; agent and hub signature fields are initialized empty. The Turkey agent was not inspected because SSH access was unavailable.

These implemented primitives support traceability but do not establish digital source authentication, immutable storage or an independently audited chain of custody. Recommended hardening includes persisted artifact hashes, verification after upload/export, signed manifests, protected audit records and separation of originals from redacted derivatives.

Field group	Reference evidence record
Identity	Tenant, site, monitor, incident, job and agent identifiers; explicit route and profile.
Time and execution	Scheduled, started, completed and received UTC times; runner and browser versions; timeout.
Network and page result	Requested/final URL, HTTP status if available, error class, duration definition and capture mode.
Artifact and retention	Artifact reference, media type, access scope and retention; agent-side SHA-256 exists, while complete persisted verification and signatures require hardening.

Implemented hashing is not a completed digital-signature system, an immutability guarantee or automatic legal admissibility.

07 / Managed distributed network model

Controlled deployment rather than anonymous volunteer capacity

The mobile deployment model uses Linux mini-PCs connected through SIM routers on selected mobile operators, placed in rented colocation facilities. Physical placement and Internet egress are separate facts: to count as a mobile path, the monitoring traffic must traverse the intended SIM connection rather than silently use the facility uplink. This architecture is owner-described, not an independent site-by-site audit.

A new path should pass staged onboarding: lawful local placement, connection-type verification, agent enrollment, version checks, clock synchronization, bounded browser execution and successful artifact upload. Credentials should be scoped per agent and revocable. Replacements and decommissioning should update inventory and invalidate old identities.

Health management separates agent degradation from a site outage. Stale heartbeats, failing browsers, backlog, clock skew and upload problems need explicit reasons. Automatic recovery is bounded and escalates repeated faults. A staged update should have a compatibility check and rollback, rather than update every path at once.

The owner reports 28 mobile points previously performed real uptime checks and browser captures. They are suspended because colocation rent, SIM Internet/data and other recurring costs require funding. Software already running on the four rented datacenter servers remains useful, but those four are hosted paths, not active mobile access. Reactivation must repeat connectivity and agent-health checks.

Recommended activation gate: fresh heartbeat, verified SIM egress, synchronized clock, successful browser capture and artifact upload; mark inconclusive observations UNKNOWN.

08 / Current network vs target expansion

Four live hosted points; 28 mobile points awaiting reactivation

Deployment status was clarified by the owner on 4 October 2026. The operating locations are Warsaw, Milan, Istanbul and London, all rented servers in datacenters. The other 28 points are Linux mini-PC/SIM-router mobile deployments, tested previously and now suspended for lack of maintenance funds. The 32-point inventory therefore does not mean 32 active points.

The owner confirms historical uptime checks and browser screenshots from the 28 mobile points. Their images were stored in R2 and automatically deleted after 10 days; those historical artifacts are no longer available for independent review. Prior testing is an owner-confirmed history, not retained proof of current availability.

Group / count	Status and interpretation
4 hosted points	Operating now: Warsaw, Milan, Istanbul and London.
28 mobile points	Previously tested; suspended pending operating funding.
80 mobile agents	First funded target: restore 28 and add 52. The four hosted points are separate.
180 target markets	Catalog plan: 1,217 entries including 114 validation slots; not active coverage.

Legacy catalog labels, including the 32-active headline and three active operator cards, must not be treated as the current deployment inventory.


09 / Network map and operator targets

A geographic planning view, not a live-coverage heat map

The map highlights countries and markets listed in the target catalog, not active coverage. Small markets may be invisible at this scale. Foundation locations are London, Warsaw, Milan and Istanbul; they do not establish full-country or mobile-operator reach. The regional totals use the directory classification: Europe 40, Africa 54, Americas 31, Asia 47 and Oceania 8. [3]

Operator targets need a maintained identity, connection type, lifecycle status and validation date. Four target slots are a planning baseline, not a claim that every country has four independent mobile operators. The live directory is the detailed reference; the map is generalized and makes no statement about disputed boundaries.

TARGET MARKETS - NOT ACTIVE COVERAGE



Foundation locations: London / Warsaw / Milan / Istanbul

Europe 40 | Africa 54 | Americas 31 | Asia 47 | Oceania 8

TARGET DIRECTORY: 180 markets | 1,217 entries including 114 validation slots. Basemap: Natural Earth, generalized 1:110m.

10 / Incident timeline and evidence lifecycle

Owner-confirmed end-to-end test; explicit confirmation and recovery scope

On 4 October 2026 the owner confirmed a successful post-fix test of the complete sequence: second DOWN, incident screenshots, received DOWN email, service recovery and received RESOLVED email. This is a first-party acceptance result, not a 100% reliability guarantee, a load test or an independently measured detection-time SLA.

Production policy should explicitly distinguish consecutive scheduled failures on one route from simultaneous results across routes. Configured thresholds, scope and recovery conditions must be recorded. The reference flow below states the two-DOWN policy; cross-route impact and late-result handling still require separately defined acceptance tests.

On the second qualifying failure, open or update one incident, enqueue capture jobs for the affected eligible routes and enqueue the DOWN notification. Evidence upload must not be a prerequisite for notifying the customer. A capture error belongs to the incident record and should not erase it. Global impact requires a separately configured cross-route policy.

01 Suspected

First qualifying DOWN; retain route result



02 Confirmed

Second DOWN; open incident and enqueue evidence + alert



03 Recovery

Configured successful checks; close scope and enqueue RESOLVED

Confirmed by the owner: DOWN -> screenshots -> received DOWN email -> recovery -> received RESOLVED email. Broader fault/retry testing remains necessary.

11 / Customer use cases and bounded value

Evidence supports decisions; it does not replace investigation

E-commerce: compare page and checkout-entry behavior across routes, collect failure screenshots and investigate whether the problem is local, regional or dependency-related. Monitoring should stop before a real purchase unless a separately authorized test environment exists. A successful landing page is not proof of a successful payment.

SaaS and enterprise portals: record login-page accessibility and approved read-only journeys. The evidence can help distinguish an identity-provider issue from a rendering or route failure. Dedicated test identities and rate limits are needed to avoid lockouts and unintended data exposure.

Hosting, CDN and WAF operations: compare challenged and unaffected paths, capture the returned screen and preserve timing and route context. An evidence package can improve an escalation conversation, but it does not automatically identify which party is responsible or guarantee acceptance of a contractual claim.

Fintech, media and public services: compare approved public or test journeys and investigate region-specific failures. Fintech checks must not initiate transfers, place orders or expose payment data. Media checks can separate an accessible shell page from blocked assets. Enterprise use adds tenant isolation, controlled exports, retention and contractual scope. Evidence supports escalation, not automatic liability attribution.

Measure value through incident reproduction rate, usable evidence completeness, confirmed notification delivery and time to a supported explanation. Revenue uplift, avoided losses and reduced response time require customer measurements; this paper supplies no invented ROI or competitive benchmark.

12 / Controlled browser scenarios

In development and testing; not deployed for customers

The owner confirms that login, cart and checkout scenarios are in development and testing, but are not yet deployed as customer features. The table describes intended boundaries, not available production journeys. Default observation remains bounded and read-only; authenticated scenarios require approved test identities, allowed actions, data scope, cleanup and cancellation rules.

Scenario	Safe boundary and interpretation
Login	Inspect public page; authenticate only with approved test credentials. Do not access unrelated customer records.
Cart	Adding an item may change state, reserve stock or trigger analytics. Use a test environment or explicitly approved disposable cart with cleanup.
Checkout	Inspect the permitted entry screen and required resources. Stop before order creation or sending sensitive data unless a sandbox is explicitly approved.
Payment	No live charge, transfer, token harvesting or control bypass. Sandbox checks do not prove that a real payment would settle.

Proposed controls: action allowlist, maximum steps, rate limits, secret redaction, emergency stop and an auditable execution trace.

13 / Third-Party Dependency Intelligence

Next phase: combine provider signals with observed impact

The proposed module lets a customer associate external providers and components with a site. Relevant categories include payments, CDN/WAF, cloud, DNS, email, SMS, identity, hosting, SaaS, development infrastructure and AI APIs. Example vendors illustrate potential adapters only; no endorsement, partnership or completed integration is implied.

A provider framework should prefer official status APIs, structured status endpoints and official feeds. Each adapter preserves source URL, external component and incident identifiers, provider update time, retrieval time, last success and parser errors. Rate limits, feed changes and licensing restrictions are part of the operating model. An unavailable feed should become unknown or stale, not operational.

Customers should select the relevant component and region where possible, rather than subscribe to every event from a global vendor. Discovery may suggest dependencies from permitted request metadata, but observed contact is not proof that a dependency is critical. Customer confirmation and clear scope prevent noisy or misleading attribution.

The module must not use captured credentials, payment data or unauthorized access to investigate a provider. Initial release should focus on reliable ingestion, component mapping, customer selection, history and alert freshness before introducing causal language. [2]

Planned capability. A provider status page is a signal, not a complete statement of customer impact.

14 / Discovery and explainable correlation

Planned phase: identify candidates without inventing certainty

Correlation compares site observations, affected routes, component dependencies and provider events within a defined time window. A candidate explanation becomes more useful when it can cite positive evidence, contradictory evidence, freshness and missing information. A confidence label should summarize a documented rule set; it is not a calibrated probability unless separately validated.

Illustrative scenario, not a measured customer incident: the site document loads, an authorized checkout read-only check fails on two independent paths, and the selected payment component reports an incident in the same period. The platform may suggest a probable dependency issue. It should also show whether unaffected routes exist, whether the component is actually used and whether the failure preceded the provider report.

Useful outputs include probable cause, likely dependency, correlated incident and possible provider issue. A categorical assertion that a vendor caused an outage requires direct supporting evidence and review. A global operational status must not override a local failure; a provider incident must not automatically imply that every customer is affected.

Evaluation should use labeled replay cases, blind reviews and negative examples: unrelated simultaneous outages, stale feeds, shared hosting paths, renamed components and sparse evidence. Rule or model versions must be preserved so an explanation can be reconstructed after an update.

Evidence is observed. Causality is inferred and must remain explicitly qualified.

15 / Incident Agent and autonomous evidence

Long-term direction with human permission boundaries

The long-term Incident Agent would assemble existing facts, request additional bounded observations and prepare an evidence package for a human responder. Its value is disciplined investigation, not unrestricted access to a customer environment. Autonomous evidence collection and autonomous remediation are different capabilities; this roadmap does not promise autonomous production repair. [2]

A reference permission model separates low-risk read-only checks, explicitly approved authenticated test journeys and prohibited production-changing actions. Purchases, money transfers, destructive edits, credential extraction, account changes and bypassing protective controls are outside default scope. Approval must identify the target, action, duration and allowed data, not merely authorize an open-ended agent.

Web content, logs and provider messages are untrusted evidence, not instructions to execute. The investigation engine needs tool allowlists, budgets, timeout and cancellation controls, secret isolation and an audit trail. Generated summaries should link back to original observations and clearly separate facts, hypotheses and recommended next steps.

Release gates should include replayable traces, adversarial content tests, permission-denial tests, emergency stop and human review of unsupported statements. Begin with evidence packaging and investigation assistance before considering broader autonomous workflows.

16 / Security, privacy and customer boundaries

Design requirements, not a compliance certification

Monitoring untrusted URLs creates risk for the service and the customer. Browser workers should be isolated from management services and other tenants, with restricted outbound access and bounded resources. Target validation must address private networks, local addresses, metadata services and redirects; initial hostname validation alone is insufficient against destination changes. These are recommended safeguards, not a verified audit of current deployment. [9]

The owner reports Google-account sign-in and customer-only views of checks and images. Code inspection confirmed ownership checks in the main result/image endpoints. A separate image-link endpoint lacked that check; it was corrected on 4 October 2026 to authorize the exact stored object through its owning check or incident, and to deny access if authorization cannot be established.

The change passed 21 authorization cases, 11 isolated HTTP endpoint cases and two public HTTPS denial checks; existing primary-image metadata was also checked without fetching customer image contents. These focused tests are not a full penetration test. R2 download links are short-lived bearer capabilities: a holder can use one until expiry, without another Google sign-in. A signed access URL is not a digital signature of the image. [12]

Screenshots and URLs may contain personal data, tokens or confidential business content. Default collection should minimize such material; query parameters, headers and bodies require careful selection and redaction. Authenticated journeys should use dedicated test identities and explicit permission. Raw secrets must never appear in public evidence packages.

The website publishes security and privacy pages. This white paper does not claim ISO certification, completed penetration testing or legal compliance approval. Before contractual enterprise use, clarify data locations, subprocessors, deletion, incident handling and any required data-processing terms with appropriate review. [4, 5]

17 / Limitations, non-claims and validation

Progress should be measurable and falsifiable

A release should be supported by repeatable tests, not only a visually successful demonstration. Controlled failures should include DNS resolution failure, invalid TLS, HTTP errors, timeouts, challenge screens, broken expected content and partial-route recovery. Test agents stopping, results arriving late, duplicate jobs, failed uploads and notification transport outages.

Operational measurements should include agent freshness, completed-job ratio, capture completeness, notification transport outcomes, backlog age, restore-test results and reconciled path inventory. Detection delay must be reported against the configured interval and confirmation rule; no absolute detection-time promise is made. Retention and deleted-domain tests must include late jobs and backup-policy review.

Physical deployment risks include supplier lead times, local contracts, SIM identity requirements, operator policy changes, shipping, power and replacement logistics. Technical risks include browser changes, WAF behavior, DNS caching, correlated infrastructure, false attribution and storage growth. Commercial risks include uncertain funding, continuing operating costs and reward obligations.

Mitigation is staged deployment, path validation, bounded automation, documented evidence semantics and transparent reporting. This paper does not guarantee uninterrupted service, full global rollout, causally correct automated conclusions or legal admissibility. A white paper explains a design and direction; it does not replace an SLA, security assessment or customer contract.

18 / Roadmap and delivery gates

A staged, multi-year program; release gates before calendar promises

The owner describes the complete platform and global network as a multi-year program. It is not realistically presented as a one-to-five-month build. The initial 80-mobile-agent target is an early funded stage, while the wider operator network, dependency intelligence, correlation and controlled Incident Agent are delivered and validated incrementally. [2]

Phase / status	Scope and proposed release gate
00 / current foundation	Availability, browser evidence, history, workspaces and DAX. Gate: repeatable DOWN/capture/recovery tests; alert delivery, deletion and path inventory reconciled.
01 / next	Provider/component registry and official-source adapters. Gate: stale/unknown handling, rate-limit tests, customer subscriptions and source provenance.
02 / planned	Dependency discovery and explainable correlation. Gate: labeled replay cases, contradiction handling, measured false attribution and evidence-linked explanations.
03 / long-term	Controlled Incident Agent and evidence packaging. Gate: bounded tools, approval controls, adversarial tests, cancellation and auditable traces.

The 12-month funded operating period is not a full-system completion deadline. No fixed year count, phase dates or 80-agent deployment date is approved; timing depends on funding, validation and operations.

19 / Commercial model and campaign context

A software service with a continuing network cost base

EyeBallMonitor is developed and operated by a two-person team. The owner-approved draft campaign goal is EUR 75,000 over 30 days. The initial mobile-network target is 80 agents in total: reactivate the 28 previously tested points and add 52 new ones. The four operating datacenter-hosted points are a separate foundation, not part of the mobile-agent total.

The owner states that this planning budget covers procurement, installation/colocation deployment and 12 months of colocation rent and SIM Internet/data for the initial mobile stage. No supplier-verified per-node prices, line-item allocation or deployment calendar have been provided. The scope is an owner-approved financing plan, not proof of funds received or an independently costed fulfillment guarantee.

Further expansion is intended to be financed by reinvesting revenue from paid customer accounts/subscriptions. Sustaining the first 80 paths also requires continuing revenue after the funded operating period. Full network completion is conditional on actual sales, operating margins, local availability and delivery capacity; neither profitability nor deployment of all 180 markets is guaranteed.

Before launch, reconcile supplier quotes, campaign deductions, logistics, contingencies and reward obligations against the stated scope. Launch dates, final reward terms and the campaign video remain separately approved. The paper offers no equity, investment return, unlimited service or guaranteed placement for an individual supporter. [6, 7]

EUR 75,000 plan: 80 mobile agents total, including 28 reactivations, plus 12 months of colocation and SIM Internet. Customer prices remain final and VAT-inclusive.

20 / FalcoSoft ITS and delivery accountability

Company identity and the responsibilities needed to scale

EyeBallMonitor is developed and operated by a two-person team at FALCOSOFT ITS EOOD, a Bulgarian company, UIC 208644390, VAT registration BG208644390. Published address: 255 Alexander Stamboliyski Blvd, Sofia, Bulgaria. Support: support@eyeballmonitor.com. Team size is owner-confirmed; no unprovided names, biographies, investment or accreditation are implied.

The team must cover product decisions, agent operations, browser/evidence engineering, customer support, deployment logistics and financial planning. Named allocation of these duties has not been supplied. An 80-agent physical network creates a support and maintenance burden; scale must follow validated operations, not only hardware purchasing.

The project should report progress through verified routes activated, accepted software milestones, operational measurements and openly stated remaining limitations. A public roadmap can explain direction; release notes should explain what actually changed. Campaign updates should distinguish procurement, installation, connectivity validation and customer availability.

The network history, two-person team, incident acceptance result and funding scope in this edition are explicit owner confirmations. Selected hashing and ownership controls were inspected in code; those checks do not constitute a comprehensive security or infrastructure audit. Future public case studies should use new authorized tests because the historical mobile captures have expired.

21 / Appendix

Terminology, architecture and evidence fields

Key terms

Agent: managed worker performing checks. Route: an observed network path, not a whole country. Evidence: artifact plus execution context. Incident: grouped failure and recovery events. Dependency: an external service used by a customer site. Correlation: related signals, not proven causality. SLO/SLA: an objective / contractual commitment requiring separately defined measurement.

Architecture and field references

See section 04 for the control-plane / agent / ingestion architecture; section 06 for the evidence-field groups and custody trail; section 10 for the confirmation / capture / recovery timeline. These are logical reference diagrams and requirements, not proprietary production schemas.

Notification and deletion controls

DOWN and RESOLVED should refer to the same incident and the same affected scope. Recovery on one route must not silently close remaining affected routes. Delivery retries need bounded backoff and an idempotency key per incident, channel, recipient and event version. Failed deliveries should be visible to operators, and recovery messages should not depend on successful delivery of the earlier DOWN email.

The product owner requires domain deletion even when monitoring is active or history exists. The reference deletion flow disables scheduling, cancels or invalidates pending work, removes domain-linked histories and evidence, and prevents late jobs from recreating deleted data. Tenant ownership must be checked before destructive actions.

Operations and capacity

Certificate renewal is an operational pipeline, not just a timer: expiry inventory, successful renewal, configuration deployment, service reload and external TLS verification. It also needs failure alerts and renewal dry runs. Automation lowers risk but cannot justify a promise that certificates or services will never fail.

Capacity planning should model checks per minute, concurrent browsers, artifact size, retention, retries and route bandwidth. For illustration only: 100 sites x 4 routes at one-minute intervals create 400 lightweight checks per minute. Browser evidence at every check would be a different and much heavier workload. Incident-triggered capture and bounded concurrency help control cost; no benchmark capacity or SLA is claimed.

References and document basis

Version 1.2, 4 October 2026, incorporates the owner's clarifications: four live hosted points, 28 suspended/tested mobile points, historical 10-day R2 deletion, the accepted incident sequence, scenarios in development, a two-person team and an 80-agent / 12-month financing plan. Focused code review covered hub controls and Poland/Milan/London agents; the image-ownership fix passed 34 security-focused tests. These are not independent deployment, security or financial audits. Product pages/catalog were checked 03-04 October; legacy labels do not override clarified status. No customer secrets or private infrastructure identifiers are reproduced.

[1] EyeballMonitor product overview

<https://eyeballmonitor.com/>

[2] EyeballMonitor public roadmap

<https://eyeballmonitor.com/roadmap>

[3] EyeballMonitor network directory

<https://eyeballmonitor.com/network>

[4] EyeballMonitor security

<https://eyeballmonitor.com/security>

[5] EyeballMonitor privacy

<https://eyeballmonitor.com/privacy>

[6] EyeballMonitor service plans

<https://eyeballmonitor.com/pricing.php>

[7] EyeballMonitor terms

<https://eyeballmonitor.com/terms>

[8] Playwright documentation: emulation

<https://playwright.dev/docs/emulation>

[9] OWASP SSRF Prevention Cheat Sheet

https://cheatsheetseries.owasp.org/cheatsheets/Server_Side_Request_Forgery_Prevention_Cheat_Sheet.html

[10] EyeballMonitor machine-readable network catalog

<https://eyeballmonitor.com/assets/data/monitoring-network.json>

[11] Natural Earth 1:110m countries basemap

https://github.com/nvkelso/natural-earth-vector/blob/master/geojson/ne_110m_admin_0_countries.geojson

[12] Cloudflare R2: presigned URLs and temporary access

<https://developers.cloudflare.com/r2/api/s3/presigned-urls/>