BOTEXORCIST DOCUMENTATION

BotExorcist Product Manual and Technical Reference

WordPress Firewall, Behavioural Traffic Intelligence and PPC Integrity

BotExorcist Lite and BotExorcist Pro

Installation, operation, evidence, PPC investigation, IP and ASN intelligence, troubleshooting, emergency recovery, and administrator procedures.

Clear Signal Intelligence (Pty) Ltd

Publication date: 16 July 2026

Contents

About this manual

This manual explains BotExorcist Lite and BotExorcist Pro, including installation, firewall operation, behavioural analysis, PPC Integrity, evidence, privacy, troubleshooting, emergency recovery, and day-to-day administration. Features that differ between Lite and Pro are identified explicitly.

---

1. Introduction

BotExorcist is a WordPress security and traffic-intelligence system developed by Clear Signal Intelligence (Pty) Ltd.

It was created to help website owners answer questions that conventional analytics and ordinary firewall tools often leave unanswered:

- Who arrived at the website?
- What did the visitor request?
- How did the visitor behave after arrival?
- Did the visitor show meaningful human interaction?
- Did the visitor probe unknown routes or suspicious parameters?
- Was the activity repeated from the same address, network, or autonomous system?
- Did the visit originate from paid advertising?
- Was the visit merely a page load, or did it show genuine interest?
- Why did the firewall observe, allow, block, or return a particular HTTP response?

BotExorcist gives the website owner an independent first-party record of traffic behaviour.

It does not assume that every technically valid browser request is a genuine person. It evaluates requests in context, records behavioural evidence, recognises human interaction, and centralises enforcement through a controlled decision layer.

The product is available in two editions:

- **BotExorcist Lite**, a free and independently functional WordPress firewall.
- **BotExorcist Pro**, a separately purchased premium product with advanced intelligence, evidence, automation, PPC analysis, reporting, and investigation capabilities.

---

2. What BotExorcist is

BotExorcist combines several related functions within one product architecture.

2.1 WordPress firewall

The firewall evaluates incoming requests, routes, parameters, behaviour, network information, and accumulated risk.

It can:

- allow normal traffic;
- observe uncertain traffic;
- recognise trusted administrators and whitelisted visitors;
- identify suspicious routes and parameters;
- record human signals;
- apply blocks when evidence supports enforcement;
- return appropriate HTTP responses;
- preserve activity records for review.

2.2 Behavioural traffic analysis

BotExorcist does not rely only on IP reputation or user-agent strings.

It can evaluate:

- scrolling;
- clicks;
- navigation;
- page depth;
- dwell behaviour;
- repeated requests;
- error patterns;
- route exploration;
- query-string behaviour;
- advertising arrival context;
- score changes over time.

2.3 IP and network intelligence

Where available, BotExorcist can associate traffic with:

- IP addresses;
- CIDR ranges;
- autonomous system numbers;
- network organisations;
- approximate geographic indicators;
- hosting providers;
- data centres;
- mobile carriers;
- repeat network behaviour.

2.4 PPC traffic intelligence

BotExorcist Pro can maintain a first-party evidence record for traffic associated with:

- Google Ads;
- Microsoft Advertising;
- Meta advertising;
- other identifiable campaign traffic.

The system helps the website owner compare paid arrivals with actual on-site behaviour.

2.5 Evidence and reporting

Depending on edition and configuration, BotExorcist can record or export:

- arrival time;
- IP and network information;
- page path;
- referrer;
- campaign context;
- behavioural events;
- human signals;
- risk-score changes;
- enforcement results;
- repeat activity;
- evidence summaries;
- CSV, spreadsheet, or PDF reports.

---

3. Why BotExorcist exists

Traditional website security products usually focus on recognisable attacks:

- repeated failed logins;
- known malicious IP addresses;
- malware signatures;
- brute-force attempts;
- known exploit paths;
- excessive request rates.

Those protections are important, but they do not describe the entire traffic problem.

A large amount of questionable activity occurs in a grey area between a clear attack and a clear human visit.

Examples include:

- silent automated arrivals;
- fake or low-quality engagement;
- paid visits with no meaningful interaction;
- repeated route enumeration;
- parameter pollution;
- scanning of nonexistent paths;
- automated browsers that imitate ordinary visitors;
- network patterns that become significant only over time;
- repeated paid arrivals from the same infrastructure.

BotExorcist was built to investigate that grey area.

The goal is not to label every unusual visitor as malicious. The goal is to collect enough evidence for a fairer and more informed decision.

---

4. Product philosophy

4.1 Behaviour over assumption

BotExorcist begins with a simple principle:

> *A valid request is not automatically a trustworthy request.*

A bot may use:

- a normal browser user agent;
- JavaScript;
- cookies;
- a residential proxy;
- a valid referrer;
- a realistic screen size.

A genuine visitor may:

- block scripts;
- browse through a corporate network;
- use accessibility technology;
- arrive through a VPN;
- avoid clicking;
- leave quickly.

For that reason, BotExorcist does not treat one signal as absolute proof.

4.2 Evidence accumulation

BotExorcist evaluates combinations of evidence.

A single unusual request may be harmless. A pattern of unknown routes, suspicious parameters, repeated errors, no human signals, and repeated activity from the same network may justify a different response.

4.3 Human redemption

A visitor who initially looks uncertain can demonstrate genuine interest through:

- scrolling;
- clicking;
- navigating;
- reaching a second page;
- spending meaningful time;
- interacting with the site.

BotExorcist can use those signals to reduce suspicion.

4.4 Independent first-party records

Advertising platforms, analytics services, hosting providers, and security systems each see only part of a visitor's journey.

BotExorcist records what occurs on the website itself.

> *We do not trust the traffic economy to investigate itself.*

This does not mean that every platform acts improperly. It means that website owners benefit from maintaining their own evidence.

---

5. BotExorcist Lite

5.1 BotExorcist Lite is a free firewall

BotExorcist Lite is a genuine, independently functional WordPress firewall.

It is not merely:

- a product launcher;
- a trial shell;
- an installer for Pro;
- an external update client;
- a catalogue of paid add-ons.

The core firewall remains functional without:

- a purchase;
- a paid licence;
- a subscription;
- a premium plugin;
- an active credit balance.

5.2 What remains free

The exact Lite feature set may evolve, but the defining rule remains:

> *The free core firewall does not stop protecting the website because optional credits are exhausted.*

Core local protection is not disabled merely because the user does not purchase BotExorcist Pro.

5.3 Lite is a product in its own right

BotExorcist Lite should be understood and presented as a complete free product.

Its purpose is to provide meaningful WordPress security and traffic visibility.

Any optional link to BotExorcist Pro is secondary.

---

6. BotExorcist Pro

BotExorcist Pro is a separately distributed premium product.

Depending on version and plan, Pro may include advanced capabilities such as:

- PPC Integrity;
- extended evidence retention;
- advanced reporting;
- evidence packs;
- complaint and dispute support;
- advanced IP and ASN intelligence;
- network cataloguing;
- query intelligence;
- deeper behavioural analysis;
- automation;
- provider-specific PPC views;
- advanced exports;
- premium support;
- additional investigation tools.

BotExorcist Pro does not change the fact that BotExorcist Lite remains a functioning free firewall.

---

7. Lite, Pro, purchases, and updates

7.1 Lite does not install Pro

BotExorcist Lite does not:

- download BotExorcist Pro;
- install BotExorcist Pro;
- activate BotExorcist Pro;
- deliver executable Pro code;
- update BotExorcist Pro;
- operate as an external plugin store.

7.2 How a user purchases Pro

The standard purchase process is:

1. The user visits the official BotExorcist website.
2. The user reviews the premium product.
3. The user completes the purchase on the website.
4. The user downloads the premium ZIP file.
5. The user uploads the ZIP manually through the WordPress dashboard.
6. The user activates the premium plugin.

7.3 How Pro updates are installed

BotExorcist Lite is not involved in Pro updates.

Where manual updating applies:

1. The user obtains the latest premium ZIP from the official website or account area.
2. The user downloads the ZIP.
3. The user uploads it manually through WordPress.
4. WordPress replaces or updates the installed premium version.

No external installer is built into BotExorcist Lite.

7.4 Lite updates

When distributed through WordPress.org, BotExorcist Lite updates are delivered through the normal WordPress.org update system.

---

8. Credits and optional services

8.1 What credits are

Credits may be used for optional services that incur processing or external-service costs.

Examples may include:

- AI-assisted analysis;
- enhanced intelligence enrichment;
- premium network lookups;
- generated investigations;
- advanced reporting;
- external processing;
- other clearly identified optional services.

8.2 What credits are not

Credits are not:

- a firewall activation fee;
- a timer that disables Lite;
- a hidden subscription requirement;
- a trial countdown;
- a licence for the free core firewall;
- a mechanism that turns off local protection.

8.3 What happens when credits reach zero

> ***The core firewall continues operating.***

When credits are exhausted, only the optional service that consumes credits may pause or become unavailable.

The local firewall must continue to perform its free core functions.

8.4 Local and external processing

The documentation and interface should clearly distinguish between:

Local functions

Functions performed within the WordPress installation, such as:

- request evaluation;
- local scoring;
- route analysis;
- whitelist checks;
- block checks;
- local logs;
- administrator protection;
- core firewall enforcement.

Optional external functions

Functions that may involve an external service, such as:

- AI analysis;
- IP enrichment;
- premium intelligence;
- generated reports;
- optional service processing.

Where data may leave the website, BotExorcist should explain:

- what is sent;
- why it is sent;
- which service receives it;
- whether the feature is optional;
- what continues working without it.

---

9. How the firewall works

A typical request may pass through the following stages:

1. A visitor requests a page, file, route, or endpoint.
2. BotExorcist checks whether the visitor or request is trusted.
3. Whitelists, administrator protections, authenticated status, and safe lanes are evaluated.
4. The requested path and query string are examined.
5. Existing IP, ASN, network, or behavioural history may be reviewed.
6. Suspicious events may increase the risk score.
7. Human interaction may reduce the risk score.
8. Observation rules determine whether more evidence is needed.
9. The Governor determines whether enforcement is justified.
10. The request is allowed, observed, blocked, redirected, or answered with an appropriate HTTP response.
11. The result is logged.
12. Relevant IP, PPC, query, network, or evidence records are updated.

The exact order may differ by version and request type, but the principle remains consistent: evidence is gathered first, and enforcement is centralised.

---

10. The Governor

The Governor is the central enforcement authority.

Other BotExorcist modules may:

- observe;
- score;
- classify;
- record;
- recommend;
- identify suspicious behaviour.

The Governor decides whether an enforcement action should occur.

10.1 Why enforcement is centralised

Centralised enforcement helps prevent:

- contradictory decisions;
- one module blocking traffic that another module trusts;
- whitelists being ignored;
- administrators being blocked by secondary logic;
- multiple systems writing conflicting rules;
- inconsistent handling across routes.

10.2 Governor principles

The Governor should respect:

- administrator protection;
- whitelists;
- safe routes;
- page-builder lanes;
- authenticated sessions;
- Observation Mode;
- confidence thresholds;
- manual overrides;
- previously approved decisions.

---

11. Observation Mode

Observation Mode allows BotExorcist to collect evidence without immediately blocking uncertain traffic.

This is important because unusual behaviour is not always malicious.

Observation Mode is especially useful during:

- initial installation;
- migration;
- proxy configuration;
- Cloudflare setup;
- page-builder testing;
- checkout testing;
- form testing;
- advertising campaign launches;
- high-traffic events;
- troubleshooting.

11.1 Advertising arrivals

Paid visitors should generally be observed carefully before enforcement.

A paid arrival that leaves quickly may be:

- a bot;
- an accidental click;
- a poor keyword match;
- a slow-loading experience;
- a genuine visitor who found the offer irrelevant;
- a privacy-conscious user;
- a person interrupted before interacting.

No single event should automatically be treated as fraud.

11.2 Recommended use

Start in Observation Mode.

Confirm:

- the correct administrator IP;
- the correct visitor IP;
- forms work;
- checkout works;
- login works;
- page builders save correctly;
- caching does not hide behaviour;
- legitimate integrations are not blocked.

Move to enforcement gradually.

---

12. Enforcement Mode

Enforcement Mode allows the Governor to apply active protection where evidence and configuration support action.

Possible enforcement outcomes may include:

- allow;
- observe;
- block;
- return HTTP 403;
- return HTTP 410;
- write a server-level block rule;
- record a manual block;
- preserve the event for later investigation.

The precise action depends on:

- request type;
- confidence;
- whitelist state;
- route;
- score;
- previous behaviour;
- administrator status;
- product configuration;
- server capabilities.

---

13. Trusted traffic and administrator protection

13.1 Administrator trust

BotExorcist should provide strong protection against accidental administrator lockout.

Trusted conditions may include:

- known administrator IP;
- successful authenticated administrator login;
- manually whitelisted IP;
- trusted ASN after successful login, where enabled;
- trusted session;
- explicit emergency bypass.

13.2 Whitelist priority

Whitelisted traffic should bypass ordinary enforcement.

Server-level rules should be ordered so that allow rules are processed before deny rules where applicable.

13.3 Trusted safe lanes

Some WordPress processes require special handling, including:

- `admin-ajax.php`;
- REST API routes;
- page-builder requests;
- editor previews;
- save operations;
- WooCommerce callbacks;
- payment-provider callbacks;
- scheduled tasks;
- trusted integrations.

Safe lanes should be specific, not broad enough to create a security bypass.

---

14. Behavioural scoring

14.1 What a score represents

A BotExorcist score is an accumulated behavioural assessment.

It is not:

- proof of identity;
- proof of criminal intent;
- proof of advertising fraud;
- a legal finding;
- an absolute classification.

It is one part of the evidence used by the system.

14.2 Risk-increasing events

Depending on version and configuration, examples may include:

- no meaningful human signal;
- honeypot interaction;
- suspicious parameter probing;
- unknown route requests;
- repeated nonexistent paths;
- suspicious redirects;
- rapid request patterns;
- repeated errors;
- reconnaissance behaviour;
- repeated network activity;
- abnormal navigation.

14.3 Risk-reducing events

Human behaviour may reduce suspicion, including:

- clicks;
- scrolling;
- navigation;
- page-two interaction;
- longer engagement;
- combined scroll and click activity;
- repeated genuine browsing.

14.4 Why scoring is cumulative

A single unusual event may be harmless.

For example:

- one 404 can be a typing mistake;
- one no-scroll session can be a quick reader;
- one data-centre IP can be a legitimate monitoring service;
- one repeated request can be a browser retry.

Patterns matter more than isolated events.

---

15. Human interaction signals

15.1 Scroll activity

Scrolling can indicate that a visitor is reading or exploring content.

Limitations:

- some pages fit entirely on screen;
- keyboard navigation may differ;
- accessibility tools may scroll differently;
- scripts may be blocked;
- some bots simulate scroll events.

15.2 Click activity

A click may indicate interaction with:

- navigation;
- buttons;
- links;
- accordions;
- tabs;
- forms;
- product controls.

A simulated click is possible, so clicks should not be treated as absolute proof.

15.3 Navigation

A visitor who moves to another page often shows stronger intent than a single page-load event.

Navigation may include:

- internal link clicks;
- menu use;
- service-page visits;
- product-page visits;
- contact-page visits;
- second-page browsing.

15.4 Dwell behaviour

Time spent on a page can help distinguish:

- immediate drive-by traffic;
- passive reading;
- delayed interaction;
- automated visits;
- genuine engagement.

Dwell time alone is not proof.

15.5 Combined signals

Combined signals are generally more meaningful than a single event.

For example:

- scroll plus click;
- click plus navigation;
- second-page activity;
- meaningful dwell plus form interaction.

15.6 Missing signals

A lack of signals is evidence, but not automatic proof of automation.

BotExorcist should record the absence of interaction without assuming guilt.

---

16. Routes, paths, and probing

16.1 Known routes

Known routes may include:

- published pages;
- posts;
- products;
- feeds;
- assets;
- login routes;
- admin routes;
- REST endpoints;
- AJAX endpoints;
- plugin routes;
- theme assets.

16.2 Unknown routes

Unknown paths may be caused by:

- old links;
- deleted content;
- typing mistakes;
- broken external links;
- search-engine remnants;
- scanning;
- exploit probing;
- route enumeration.

16.3 Route probing

Repeated requests for known exploit targets, common plugin paths, backup files, environment files, or nonexistent administrative routes may indicate reconnaissance.

Examples include repeated attempts to find:

- configuration files;
- backup archives;
- vulnerable plugin paths;
- old admin routes;
- hidden panels;
- development files;
- exposed logs.

16.4 Why one unknown route is not enough

One unknown path is often harmless.

Repeated, structured, or rapidly changing unknown paths may be more significant.

---

17. Query and parameter intelligence

Query parameters can be legitimate, suspicious, or ambiguous.

17.1 Legitimate parameters

Examples include:

- campaign tracking;
- analytics identifiers;
- search terms;
- pagination;
- filters;
- product selections;
- language settings;
- form tokens;
- referral tags.

17.2 Suspicious parameters

Suspicious behaviour may include:

- probing for executable parameters;
- unexpected file paths;
- injection patterns;
- repeated random keys;
- malformed values;
- parameter pollution;
- attempts to alter WordPress behaviour;
- repeated unknown parameters.

17.3 Query Intelligence decisions

A query-intelligence system may classify or recommend:

- block;
- do not block;
- observe;
- manually review;
- preserve history;
- apply override.

17.4 Administrator exclusion

Administrator requests should be excluded from automated query blocking where appropriate.

17.5 Historical decisions

A good query-intelligence system preserves:

- prior observations;

- manual decisions;
- overrides;
- previous verdicts;
- change history.

---

18. HTTP responses and enforcement actions

18.1 Allow

The request is permitted.

18.2 Observe

The request is permitted while evidence is recorded.

18.3 HTTP 403 Forbidden

A 403 response indicates that access is understood but refused.

It may be appropriate for:

- blocked visitors;
- denied routes;
- server-level restrictions;
- disallowed activity.

18.4 HTTP 404 Not Found

A 404 indicates that the requested resource was not found.

This is appropriate for ordinary missing content where the resource may have never existed or where normal not-found handling is desired.

18.5 HTTP 410 Gone

A 410 indicates that a resource is intentionally unavailable and should be treated as gone.

It may be used for:

- known spam routes;
- junk parameter URLs;
- permanently removed polluted URLs;
- automated probing paths;
- unwanted search-index remnants.

18.6 Search-engine implications

A 410 can encourage search engines to remove a URL more quickly than a normal 404 in some circumstances.

It should not be used indiscriminately.

Real pages, assets, editors, previews, and valid routes must be protected from accidental 410 responses.

---

19. Whitelists, blacklists, IP ranges, and CIDR

19.1 Whitelists

A whitelist explicitly trusts:

- an IP address;
- a CIDR range;
- a known administrator;
- a trusted network;
- another approved source.

Whitelists should override ordinary blocking logic.

19.2 Blacklists

A blacklist explicitly blocks:

- an IP address;
- a network range;
- a known hostile source;
- a manually confirmed threat.

19.3 CIDR ranges

CIDR notation describes a network range.

Examples:

- `192.0.2.0/24`
- `2001:db8::/32`

Blocking an entire range is more powerful and more dangerous than blocking one IP.

19.4 Caution with network-wide blocks

Do not block an entire ASN or CIDR range merely because one address behaved suspiciously.

Large ranges may include:

- mobile users;
- business customers;
- shared hosting;
- cloud infrastructure;
- legitimate services;
- search crawlers;
- payment providers.

Network-wide action should require strong evidence.

---

20. IP intelligence

20.1 What an IP address can indicate

An IP address may help identify:

- approximate network origin;
- service provider;
- hosting provider;
- mobile carrier;
- data centre;
- repeated activity;
- associated ASN;
- country or city estimate.

20.2 What an IP address cannot prove

An IP address does not necessarily identify:

- a natural person;
- a household member;
- an employee;
- a device owner;
- the person who clicked an ad;
- the person responsible for an action.

20.3 Shared addresses

One IP may represent:

- an office;
- a school;
- a mobile network;
- a hotel;
- a VPN;
- a public proxy;
- a carrier-grade NAT network;
- a cloud service;
- many unrelated users.

20.4 Geolocation limitations

IP geolocation is approximate.

City-level results may be:

- outdated;
- based on provider registration;
- mapped to a network hub;
- mapped to an ISP office;
- inaccurate for mobile traffic;
- inaccurate for VPNs.

21. ASN and network intelligence

21.1 What an ASN is

An autonomous system number identifies a network or group of networks operated under a common routing policy.

21.2 Why ASNs matter

ASN information can help identify:

- repeat traffic from the same network;
- hosting providers;
- cloud networks;
- mobile carriers;
- corporate infrastructure;
- data-centre traffic;
- network-level patterns.

21.3 Repeat ASN activity

Repeated suspicious events across different IP addresses within the same ASN may reveal a broader pattern.

However, the pattern must be interpreted carefully.

A large ASN may serve millions of unrelated users.

21.4 ASN complaints

Where appropriate, abuse reports may be sent to a network provider.

Reports should include:

- factual event details;
- timestamps;
- affected paths;
- source IPs;
- relevant evidence;
- previous correspondence;
- a clear request.

22. IP Intelligence Map and Catalogue

The IP Intelligence Map and Catalogue may provide:

- recently observed IPs;
- threat indicators;
- approximate location;
- ASN;
- organisation;
- network information;
- historical activity;
- manual notes;
- enrichment;
- abuse-report options.

22.1 Map interpretation

The map is a visual aid.

It should not be treated as exact proof of a visitor's physical location.

22.2 Threat indicators

A visual threat indicator may represent:

- a new suspicious IP;
- recent activity;
- a block;
- a score threshold;
- a manually classified source.

The interface should explain exactly what each indicator means.

---

23. PPC Integrity

PPC Integrity is a first-party advertising-traffic intelligence system.

It helps website owners compare paid traffic with on-site behaviour.

23.1 Supported sources

Depending on version and configuration, PPC Integrity may recognise:

- Google Ads;
- Microsoft Advertising;
- Meta advertising;
- tagged campaign traffic;
- identifiable paid-arrival parameters.

23.2 What PPC Integrity records

Possible fields include:

- arrival timestamp;
- landing page;
- referrer;
- provider;
- campaign parameters;
- IP;
- ASN;
- organisation;
- approximate location;
- device context;
- dwell;
- scroll;
- clicks;
- navigation;
- repeat activity;
- score;
- action.

23.3 Click versus interest

A click is not the same as a genuine prospect.

A paid click may lead to:

- meaningful engagement;
- immediate exit;
- accidental arrival;
- bot activity;
- poor keyword relevance;
- slow-page abandonment;
- competitor research;
- invalid traffic;
- legitimate but low-intent browsing.

PPC Integrity helps document the difference between a charged arrival and observable on-site interest.

23.4 Drive-by activity

A drive-by visit may involve:

- an arrival;
- little or no dwell;
- no scroll;
- no click;
- no navigation;
- immediate departure.

Drive-by behaviour is evidence of low engagement, not automatic proof of fraud.

23.5 Repeat IP and ASN patterns

Repeated arrivals from the same IP or network may be significant when combined with:

- repeated no-signal behaviour;
- repeated paid arrivals;
- identical landing patterns;
- unusual timing;
- repeated campaign identifiers;
- suspicious routes;
- high-frequency activity.

23.6 Provider views

PPC Integrity may separate traffic by provider to support clearer reporting and investigation.

---

24. Advertising evidence and interpretation

24.1 Evidence is not a verdict

BotExorcist records evidence and patterns.

It does not automatically establish:

- criminal fraud;
- platform wrongdoing;
- advertiser liability;
- contractual breach;
- entitlement to a refund.

24.2 Stronger evidence

Evidence is generally stronger when several independent indicators align.

Examples:

- paid arrival;
- repeated IP or ASN;
- no human signals;
- repeated campaign pattern;
- suspicious route activity;
- repeated short sessions;
- consistent timing;
- identical behaviour over multiple visits.

24.3 Responsible language

Preferred language includes:

- suspicious activity;
- low-signal traffic;
- non-human indicators;
- bot-like behaviour;
- repeated network pattern;
- activity requiring review;
- evidence consistent with automation.

Avoid unsupported claims such as:

- proven fraud;
- criminal click;
- guaranteed bot;
- deliberate theft.

---

25. Evidence Ledger

The Evidence Ledger is intended to preserve a structured record of relevant events.

Possible fields include:

Field	Meaning
Timestamp	When the event occurred
IP address	Source address observed by the site
CIDR	Associated network range
ASN	Associated autonomous system
Organisation	Network or registered organisation
Country	Approximate country
City	Approximate city
Path	Requested route
Query	Query string or parameter context
Referrer	Reported referring source
Provider	Advertising provider where identified
Campaign	Campaign identifier where available
Event	Behavioural or security event
Human signals	Scroll, click, navigation, or related interaction
Score before	Risk score before the event
Score after	Risk score after the event
Action	Allow, observe, block, or other outcome
Block state	Whether a block existed or was applied
Source	Automatic, manual, administrator, or system source
Dwell	Time-related engagement value
Scroll	Scroll activity
Clicks	Click activity
Notes	Manual or system notes

Not every field will be available for every event.

26. Reports and exports

Depending on edition, BotExorcist may support:

- CSV exports;
- Excel-compatible exports;
- PDF evidence packs;
- date-range reports;
- provider reports;
- campaign reports;
- ASN reports;
- IP reports;
- filtered evidence;
- complaint attachments.

26.1 Preserve raw evidence

When preparing an investigation:

- keep the original export;
- create a working copy;
- avoid altering timestamps;
- record the export date;
- preserve relevant server logs;
- preserve campaign records;
- note any manual classifications.

---

27. Complaint and dispute support

BotExorcist may assist with preparing factual communications to:

- advertising providers;
- hosting providers;
- network operators;
- abuse desks;
- agencies;
- clients.

27.1 What a complaint should contain

A useful complaint should include:

- a clear subject;
- date range;
- affected campaign or site;
- source IPs;
- relevant ASNs;
- timestamps;
- request paths;
- behavioural summary;
- supporting export;
- the action requested.

27.2 Cooldowns and escalation

Repeated complaints should not be sent too frequently.

A complaint cooldown can help prevent:

- duplicate messages;
- abuse-desk fatigue;
- inconsistent evidence windows;
- unnecessary escalation.

27.3 Legal caution

BotExorcist is not a law firm.

Generated or assisted correspondence should be reviewed by the user before submission.

---

28. Login Shield

Login Shield can help protect WordPress authentication.

Possible features include:

- trust after successful login;
- optional IP-based login controls;
- optional ASN-based controls;
- failed-login observation;
- session trust;
- administrator safety;
- trusted network recording.

28.1 Mobile and changing networks

Strict IP locking can create problems when a user:

- moves between Wi-Fi and mobile data;
- uses a VPN;
- changes offices;
- uses a dynamic IP;
- logs in through a mobile carrier.

For this reason, strict login locking should be optional and clearly explained.

28.2 Recovery

Users must know how to recover access through:

- hosting file manager;
- FTP;
- SSH;
- `.htaccess`;
- plugin-folder renaming.

See Emergency recovery: blocked or locked out.

---

29. WordPress compatibility

BotExorcist should account for normal WordPress processes, including:

- public page requests;
- admin requests;
- `admin-ajax.php`;
- REST API;
- WP-Cron;
- login;
- registration;
- media;
- feeds;
- preview requests;
- heartbeat;
- plugin callbacks;
- theme requests;
- WooCommerce;
- forms;
- membership systems.

Safe handling should be specific and auditable.

---

30. Page-builder compatibility

Page builders often use unusual request patterns.

BotExorcist may require safe lanes for:

- Divi;
- Elementor;
- Gutenberg;
- preview frames;
- AJAX saves;
- REST requests;
- dynamic CSS generation;
- editor assets;
- template saves;
- frontend editing.

If a page builder cannot save:

1. Switch to Observation Mode.
2. Confirm administrator trust.
3. Review the blocked request.
4. Identify the exact route.
5. Add a narrow safe lane.
6. Retest.
7. Avoid broad exclusions.

---

31. Hosting, CDN, proxy, and caching environments

31.1 Real visitor IP

A proxy or CDN may cause WordPress to see the proxy IP instead of the visitor IP.

Examples include:

- Cloudflare;
- reverse proxies;
- load balancers;
- managed hosting;
- web application firewalls;
- caching networks.

BotExorcist must be configured to trust only valid forwarding headers from trusted proxies.

31.2 Misconfigured headers

Trusting arbitrary `X-Forwarded-For` headers can allow spoofing.

Real-IP detection should use:

- known proxy ranges;
- validated headers;
- hosting-provider guidance;
- controlled configuration.

31.3 Caching

Caching may affect:

- JavaScript signals;
- page variations;
- event recording;
- logged-in behaviour;
- administrator bypasses;
- cache-clearing after updates.

---

32. Performance and resource usage

Security logging can generate significant data on busy websites.

Performance considerations include:

- request-time processing;
- database writes;
- log retention;
- caching;
- cleanup schedules;
- IP enrichment;
- external lookups;
- report generation;
- map data;
- high-volume PPC traffic.

Recommended practices:

- use sensible retention periods;
- avoid unnecessary external lookups;
- archive important reports;
- review database growth;
- test on staging;
- monitor slow queries;
- schedule cleanup outside peak periods.

---

33. Data storage and retention

BotExorcist may store:

- settings;
- activity logs;
- IP records;
- ASN records;
- PPC records;
- evidence records;
- query decisions;
- login trust records;
- complaint history;
- manual notes.

33.1 Retention

Retention should be appropriate to:

- site traffic;
- investigation needs;
- privacy obligations;
- server capacity;
- legal requirements;
- advertising dispute windows.

33.2 Reset functions

A reset function should explain exactly what it removes.

Some dashboards may retain:

- IP records;
- evidence history;
- PPC summaries;
- manual classifications.

Never assume that one reset clears every data category.

33.3 Uninstallation

Documentation should state whether uninstallation removes:

- settings;
- logs;
- custom tables;
- server rules;
- cached data;
- licence data.

Users should export important evidence before deletion.

---

34. Privacy and data handling

BotExorcist processes security and traffic information.

Depending on configuration, this may include personal data such as IP addresses.

34.1 Data minimisation

Collect only what is necessary for:

- security;
- traffic analysis;
- evidence;
- troubleshooting;
- support;
- product operation.

34.2 Website-owner responsibility

The website owner is responsible for:

- updating the privacy policy;
- choosing retention periods;
- identifying lawful bases;
- controlling access to logs;
- responding to data requests;
- protecting exported reports;
- understanding local law.

34.3 External services

Where an optional feature sends data to an external service, the product should disclose:

- the data sent;
- the recipient;
- the purpose;
- whether the feature is optional;
- the applicable privacy terms.

---

35. Security architecture

BotExorcist should follow secure WordPress development practices, including:

- capability checks;
- nonce verification;
- input validation;
- output escaping;
- prepared database queries;
- secure file handling;
- controlled remote requests;
- least privilege;
- administrator protection;
- fail-safe defaults;
- explicit update boundaries.

35.1 Remote code

BotExorcist Lite must not download or execute external plugin code.

Optional remote services may return data or analysis, but not executable plugin updates.

---

36. Installation

36.1 Before installation

Before activating BotExorcist:

- back up the website;
- confirm hosting control-panel access;
- confirm FTP or SSH access;
- record the current administrator IP;
- know where the WordPress plugin directory is;
- know where `.htaccess` is stored;
- confirm how to disable a plugin through file access;
- begin in Observation Mode.

36.2 Installing Lite through WordPress

1. Open the WordPress dashboard.
2. Go to **Plugins -> Add New**.
3. Search for **BotExorcist Lite**.
4. Select the official plugin.
5. Click **Install Now**.
6. Click **Activate**.
7. Complete initial configuration.
8. Confirm administrator trust.
9. Remain in Observation Mode while testing.

36.3 Manual Lite installation

1. Download the official Lite ZIP.
2. Open **Plugins -> Add New -> Upload Plugin**.
3. Select the ZIP.
4. Install.
5. Activate.
6. Confirm safe access.

36.4 Installing Pro manually

1. Purchase BotExorcist Pro from the official website.
2. Download the Pro ZIP.
3. Open **Plugins -> Add New -> Upload Plugin**.
4. Select the ZIP.
5. Install.
6. Activate.
7. Complete licence or entitlement steps where applicable.
8. Confirm that protection and settings are correct.

36.5 Important purchase and installation notice

> ***Emergency access recovery***
> *Before activating enforcement, make sure you can access your hosting file manager, FTP, or SSH.*
> *If you are ever blocked or locked out:*
>
> 1. *Open `.htaccess` and remove the rule blocking your IP.*
> 2. *Go to `/wp-content/plugins/`.*
> 3. *Rename `botexorcist` to `botexorcist.off`.*
> 4. *Log in to WordPress normally.*
> 5. *Rename the folder back to `botexorcist`.*
> 6. *Reactivate the plugin if required.*
> 7. *Whitelist your IP and confirm correct IP detection before restoring enforcement.*
>
> *For Lite, rename `botexorcist-lite` to `botexorcist-lite.off`.*

This notice should also appear on:

- the purchase confirmation page;
- the download email;
- the installation guide;
- the ZIP README;
- the built-in Help page.

---

37. Initial configuration

37.1 Record the administrator IP

Confirm that BotExorcist sees the same public IP that your browser uses.

If Cloudflare or a proxy is present, verify forwarded-IP configuration.

37.2 Begin in Observation Mode

Do not activate strict enforcement immediately on an unfamiliar environment.

37.3 Test critical functions

Test:

- login;
- logout;
- password reset;
- forms;
- checkout;
- account pages;
- page-builder saves;
- REST requests;
- AJAX actions;
- scheduled tasks;
- payment callbacks;
- webhooks.

37.4 Review trusted routes

Do not whitelist entire systems unnecessarily.

Add the narrowest safe rule that solves the problem.

---

38. Recommended operating procedure

38.1 Day one

- install;
- confirm backups;
- verify administrator IP;
- enable Observation Mode;
- test forms and editors;
- inspect initial events;
- verify real visitor IP;
- confirm server-rule access.

38.2 First week

- review unknown routes;
- review suspicious parameters;
- identify legitimate integrations;
- inspect repeated IPs and ASNs;
- review paid traffic;
- add necessary whitelists;
- investigate false positives;
- avoid overblocking.

38.3 Weekly

- review new blocks;
- review new network patterns;
- inspect PPC anomalies;
- confirm administrator trust;
- export significant evidence;
- clear incorrect classifications;
- review storage growth.

38.4 Monthly

- compare PPC periods;
- review repeat networks;
- update trusted rules;
- archive reports;
- review retention;
- check version and compatibility;
- test emergency access.

39. Understanding the dashboard

Dashboard metrics should always be interpreted in context.

39.1 Unique IPs

The number of distinct observed addresses.

One person may use several IPs, and many people may share one IP.

39.2 Bot behaviour

Events or sessions that contain signals associated with automation or suspicious activity.

This is not always equivalent to confirmed malicious bots.

39.3 Human signals

Recorded interaction such as:

- scroll;
- click;
- navigation;
- page depth;
- dwell.

39.4 Honeypot activity

Interaction with a hidden or deceptive element intended to attract automated behaviour.

Honeypot interaction is generally strong evidence but should still be logged and reviewed.

39.5 Charged without interest

Paid arrivals that show little or no meaningful engagement.

This is a behavioural category, not an automatic fraud verdict.

39.6 Repeat ASNs

Networks that appear repeatedly in relevant traffic.

Large networks require cautious interpretation.

---

40. False positives

False positives can occur because legitimate users may look unusual.

Common causes include:

- shared networks;
- mobile carriers;
- VPNs;
- corporate proxies;
- script blockers;
- accessibility tools;
- privacy browsers;
- aggressive caching;
- monitoring services;
- SEO crawlers;
- payment callbacks;
- webhook systems;
- page builders.

40.1 Investigation process

When investigating a possible false positive:

1. Identify the exact event.
2. Check whether the request was authenticated.
3. Check the route and parameter.
4. Check prior history.
5. Check IP and ASN.
6. Check human signals.
7. Check whether the route belongs to a plugin or service.
8. Temporarily observe rather than block.
9. Add the narrowest safe exception.
10. Retest.

41. Emergency recovery: blocked or locked out

This section should be easy to find from search engines, purchase emails, README files, and support pages.

41.1 Use this procedure when

Use this procedure if:

- your IP was blocked;
- WordPress login is inaccessible;
- the dashboard returns HTTP 403;
- the website owner is locked out;
- a server-level rule continues to block access;
- BotExorcist prevents normal administrator login.

41.2 Before making changes

Use one of the following:

- hosting file manager;
- FTP;
- SFTP;
- SSH.

Create backups of:

- `.htaccess`;
- the plugin folder;
- relevant server configuration.

41.3 Remove your IP from `.htaccess`

Open the `.htaccess` file in the WordPress root directory.

Locate the BotExorcist-managed block section.

Remove only the rule that blocks your current IP address.

Do not delete unrelated WordPress rewrite rules.

Save the file.

41.4 Temporarily disable BotExorcist

Open:

```
/wp-content/plugins/
```

For BotExorcist Pro, rename:

```
botexorcist
```

to:

```
botexorcist.off
```

For BotExorcist Lite, rename:

```
botexorcist-lite
```

to:

```
botexorcist-lite.off
```

WordPress can no longer load the plugin while the folder has the temporary name.

41.5 Log in normally

Open:

```
https://yourdomain.example/wp-admin/
```

Log in with the normal administrator credentials.

41.6 Restore the plugin folder

Rename:

```
botexorcist.off
```

back to:

```
botexorcist
```

or:

```
botexorcist-lite.off
```

back to:

```
botexorcist-lite
```

41.7 Reactivate if necessary

WordPress may mark the plugin inactive after the folder was unavailable.

Open **Plugins** and reactivate BotExorcist.

41.8 Correct the cause before restoring enforcement

Before returning to enforcement:

- add the current administrator IP to the whitelist;
- confirm correct visitor-IP detection;
- review Cloudflare or proxy configuration;
- review the event that caused the block;
- switch to Observation Mode;
- test login;
- test logout;
- test the dashboard;
- confirm `.htaccess` rule order.

41.9 Important note

Renaming the plugin folder:

- temporarily disables the plugin;
- does not normally delete settings;
- does not normally delete logs;
- does not remove database records;
- may leave server-level rules active until those rules are removed.

That is why `.htaccess` must also be checked.

---

42. Troubleshooting

42.1 Administrator blocked repeatedly

Check:

- current public IP;
- dynamic IP changes;
- proxy headers;
- VPN;
- mobile network;
- whitelist;
- trusted login;
- ASN locking;
- `.htaccess` order.

42.2 Page builder cannot save

Check:

- administrator trust;
- AJAX;
- REST;
- preview route;
- nonce;
- editor safe lane;
- request log;
- server cache.

42.3 Checkout or forms fail

Check:

- form endpoint;
- AJAX request;
- REST request;
- payment callback;
- webhook;
- anti-bot scripts;
- session cookie;
- server firewall.

42.4 IP appears incorrect

Check:

- Cloudflare;
- reverse proxy;
- hosting headers;
- trusted proxy ranges;
- spoofable forwarded headers;
- cached values.

42.5 Logs are empty

Check:

- plugin activation;
- database tables;
- caching;
- write permissions;
- cleanup settings;
- filters;
- date range;
- administrator exclusions.

42.6 PPC arrivals are missing

Check:

- campaign parameters;
- referrer;
- landing-page cache;
- JavaScript;
- provider detection;
- date range;
- consent tools;
- redirects.

42.7 Credits do not update

Check:

- account status;
- network access;
- service endpoint;
- server time;
- licence association;
- cache;
- optional-service configuration.

Remember: the core firewall must continue operating even when optional credits are unavailable.

42.8 HTTP 410 appears on a real page

Check:

- route recognition;
- page-builder preview routes;
- query rules;
- permalink cache;
- custom post types;
- manual 410 rules;
- stored query decisions;
- CDN cache.

Switch to Observation Mode before retesting.

42.9 Administrator AJAX or REST requests are blocked

Check:

- authenticated status;
- nonce verification;
- current administrator IP;
- `admin-ajax.php` safe lane;
- REST safe lane;
- request method;
- plugin-specific routes.

42.10 Server cache shows old behaviour

Clear:

- WordPress cache;
- object cache;
- Redis or Memcached;
- server cache;
- CDN cache;
- browser cache.

Then verify the active plugin version and request again.

---

43. Updates and version management

43.1 BotExorcist Lite updates

When listed on WordPress.org, Lite updates are delivered through the normal WordPress.org update mechanism.

BotExorcist Lite does not obtain executable plugin updates from BotExorcist.com.

43.2 BotExorcist Pro updates

Where manual updating applies:

1. Download the current premium ZIP from the official source.
2. Back up the website.
3. Upload the ZIP through WordPress.
4. Confirm the replacement or update.
5. Test critical functions.
6. Confirm the installed version.

BotExorcist Lite is not involved in this process.

43.3 Backups before updating

Before updating:

- back up files;
- back up the database;
- export important evidence;
- record the current version;
- record custom rules;
- test on staging where possible.

43.4 Database migrations

An update may modify or create:

- tables;
- indexes;
- options;
- stored settings;
- evidence structures;
- cached records.

Migration routines should be safe, versioned, and recoverable.

43.5 Rollback

A rollback plan should preserve:

- database compatibility;
- rule consistency;
- server-level configuration;

- evidence history;
- licence state.

43.6 Version-aware documentation

Documentation links may include the installed version so users can see guidance relevant to their build.

Example:

```
https://botexorcist.com/documentation/?version=2.0.284#observation-mode
```

---

44. Licensing

44.1 Lite licensing

BotExorcist Lite is free.

Its core firewall does not require a paid entitlement.

44.2 Pro licensing

BotExorcist Pro may require:

- purchase;
- licence key;
- site entitlement;
- subscription;
- renewal;
- support entitlement.

The applicable commercial terms should be stated clearly before purchase.

44.3 Expired licence

Documentation should explain exactly what happens when a Pro licence expires.

A licence expiry should not silently disable unrelated free firewall functionality.

Possible consequences may include:

- loss of premium support;
- loss of access to premium downloads;
- loss of selected premium services;
- inability to receive premium updates.

44.4 Staging and development

Licence rules should explain:

- staging domains;
- localhost;
- development sites;
- domain changes;
- migrations;
- agency use;
- multisite use.

44.5 Licence privacy

Licence requests should transmit only the information required to:

- validate entitlement;
- associate a site;
- provide service access;
- manage support or updates.

45. What BotExorcist does not claim

BotExorcist does not claim to:

- identify the natural person behind every IP;
- conclusively prove click fraud on its own;
- guarantee advertising refunds;
- establish legal liability;
- block every malicious request;
- replace website backups;
- replace malware scanning;
- replace secure hosting;
- guarantee exact geolocation;
- guarantee that every low-signal visitor is automated;
- guarantee that every data-centre visitor is malicious;
- eliminate all false positives;
- replace human judgement.

A responsible security product explains its limits.

---

46. BotExorcist and other security layers

BotExorcist can operate as part of a layered security strategy.

46.1 Hosting firewall

A hosting firewall can block traffic before WordPress loads.

BotExorcist adds application-level behavioural and WordPress-specific context.

46.2 Cloudflare

Cloudflare can provide:

- CDN;
- DDoS protection;
- caching;
- bot controls;
- firewall rules.

BotExorcist provides first-party WordPress and on-site behavioural evidence.

46.3 Malware scanner

A malware scanner looks for infected or changed files.

BotExorcist primarily evaluates traffic and behaviour.

46.4 Analytics

Analytics platforms summarise visitor activity.

BotExorcist records security and evidence-oriented events that analytics tools may not preserve.

46.5 Advertising invalid-click systems

Advertising platforms use their own systems.

BotExorcist provides an independent website-owner record.

46.6 Server logs

Server logs show requests and responses.

BotExorcist adds WordPress context, behavioural signals, scoring, and administrative interpretation.

---

47. Use cases

47.1 Small business website

Protect a local-service website from:

- probes;
- spam routes;
- suspicious parameters;
- repeated automated activity.

47.2 PPC landing page

Compare paid arrivals with:

- engagement;
- repeat networks;
- no-signal visits;
- campaign behaviour.

47.3 WooCommerce shop

Monitor:

- login;
- checkout;
- product browsing;
- bots;
- scraping;
- suspicious account activity.

47.4 Agency

Use BotExorcist to provide:

- client traffic reviews;
- evidence exports;
- PPC investigations;
- network analysis;
- operational reports.

47.5 Publisher

Review:

- scraping;
- content extraction;
- repeated automated sessions;
- route probing.

47.6 Membership site

Protect:

- login;
- registration;
- account areas;
- repeated failed access;
- suspicious sessions.

47.7 Lead-generation website

Assess:

- paid arrivals;
- form engagement;
- contact-page navigation;
- repeated low-signal traffic;
- suspicious campaign patterns.

47.8 High-risk login area

Combine:

- trusted sessions;
- administrator protection;
- failed-login analysis;
- optional IP or ASN trust;
- emergency recovery procedures.

47.9 Site under parameter spam

Use query intelligence and 410 handling to:

- identify junk URLs;
- preserve real routes;
- reduce index pollution;
- observe repeat sources.

---

48. Example investigations

48.1 Genuine paid visitor

Observed behaviour:

- arrives through paid campaign;
- spends time on landing page;
- scrolls;
- clicks;
- visits another page;
- views contact information.

Interpretation:

- shows meaningful human interest;
- risk should be reduced;
- no block is justified.

48.2 Drive-by paid arrival

Observed behaviour:

- paid arrival;
- no scroll;
- no click;
- very short dwell;
- no second page.

Interpretation:

- low engagement;
- preserve as evidence;
- not automatically fraud.

48.3 Automated probe

Observed behaviour:

- several nonexistent WordPress paths;
- suspicious parameters;
- rapid sequence;
- no human signals;
- repeated errors.

Interpretation:

- strong automation or reconnaissance indicators;
- enforcement may be justified.

48.4 False positive avoided

Observed behaviour:

- unusual shared network;

- first request appears suspicious;
- visitor scrolls;
- clicks;
- navigates;
- completes a form.

Interpretation:

- human redemption prevents premature blocking.

48.5 Repeat network PPC pattern

Observed behaviour:

- several paid arrivals;
- changing IPs within one network;
- repeated short sessions;
- no interaction;
- same landing page;
- repeated campaign context.

Interpretation:

- preserve evidence;
- compare time patterns;
- review network ownership;
- avoid claiming fraud without supporting evidence.

48.6 Parameter-pollution campaign

Observed behaviour:

- repeated junk parameters;
- URLs appearing in search indexes;
- no legitimate page variation;
- repeat source networks.

Interpretation:

- classify parameters;
- protect legitimate tracking parameters;
- return appropriate 410 responses where justified;
- monitor search-index removal.

---

49. Built-in Help and Documentation

BotExorcist should provide both local help and links to the complete online manual.

49.1 Help page inside the plugin

Recommended menu item:

BotExorcist -> Help & Documentation

Recommended layout:

Search

A local help search box.

Getting started

- Install safely
- Begin in Observation Mode
- Confirm administrator IP
- Review the first traffic records

Common tasks

- Whitelist an IP
- Remove a block
- Review PPC traffic
- Export evidence
- Investigate a false positive

Understand BotExorcist

- Behavioural scoring
- Human signals
- IP intelligence
- ASN intelligence
- Query intelligence
- The Governor

Emergency recovery

- Blocked or locked out
- Rename the plugin folder
- Remove `.htaccess` block
- Restore access
- Whitelist the administrator

Free firewall, credits, and Pro

- What remains free
- What credits are used for
- What happens at zero credits
- How Pro is purchased
- How Pro is updated

Support

- Export diagnostics
- Contact support
- Open complete documentation

49.2 Contextual help links

Each complex setting should link directly to the correct documentation anchor.

Examples:

- Observation Mode -> `#11-observation-mode`
- Credits -> `#8-credits-and-optional-services`
- HTTP 410 -> `#18-http-responses-and-enforcement-actions`
- ASN -> `#21-asn-and-network-intelligence`
- Evidence Ledger -> `#25-evidence-ledger`
- Locked out -> `#41-emergency-recovery-blocked-or-locked-out`

49.3 Essential offline help

The plugin should keep a local copy of emergency instructions because a locked-out user may not be able to access the online manual.

49.4 WordPress screen Help tabs

Each BotExorcist administration page should use the standard WordPress **Help** tab where practical.

Examples:

PPC Integrity screen

- Understanding summary cards
- Interpreting low-signal visits
- Evidence limitations
- Exporting reports

Query Intelligence screen

- Suspicious parameters
- Block versus Do Not Block
- Overrides
- Historical decisions

Activity Log screen

- Event meanings
- Filters
- Score changes
- Block state

49.5 Search behaviour

The documentation search should index:

- headings;
- settings;
- error messages;

- glossary terms;
- HTTP statuses;
- troubleshooting entries;
- dashboard labels;
- feature names.

---

50. Frequently asked questions

Is BotExorcist Lite really free?

Yes. BotExorcist Lite is a free, independently functional WordPress firewall.

Does the firewall stop when credits reach zero?

No. The core firewall continues operating. Credits apply only to identified optional services.

Is Lite a trial?

No. Lite is not a countdown trial that disables the core firewall.

Does Lite install Pro?

No. Lite does not download, install, activate, deliver, or update Pro.

How do I install Pro?

Purchase Pro, download the ZIP, and upload it manually through WordPress.

How do Pro updates work?

Where manual updates apply, download the new ZIP and upload it manually through WordPress.

Can BotExorcist identify a person from an IP?

No. An IP address does not necessarily identify a natural person.

Does low engagement prove fraud?

No. It is evidence of low engagement, not automatic proof of fraud.

Can BotExorcist guarantee an advertising refund?

No.

Does BotExorcist replace Cloudflare?

No. The products address different layers and may be used together.

Can BotExorcist block administrators?

It is designed to protect administrators, but configuration or IP-detection problems can cause lockout. Emergency recovery instructions are provided.

What should I do if I am locked out?

Remove the blocking rule from `.htaccess`, rename the plugin folder, log in, restore the folder, reactivate if necessary, whitelist your IP, and verify IP detection.

Why start in Observation Mode?

It allows BotExorcist to learn the environment and helps prevent false positives during setup.

Can a shared IP cause confusion?

Yes. Offices, mobile networks, hotels, VPNs, and proxies may serve many users through one address.

Is geolocation exact?

No. IP geolocation is approximate.

Does BotExorcist replace malware scanning?

No.

Does BotExorcist work with page builders?

It is designed to support major builders, but safe lanes may be needed for specific requests.

Does BotExorcist work with WooCommerce?

It can, provided checkout, AJAX, REST, payment, and webhook routes are tested and correctly handled.

What is a honeypot?

A deceptive element intended to attract automated activity.

What is an ASN?

A number identifying a network or routing organisation.

What is CIDR?

A notation used to describe an IP network range.

Why return 410?

A 410 can indicate that an unwanted or removed URL is permanently gone.

Does one 404 mean a bot?

No.

Can bots simulate clicks?

Yes. That is why BotExorcist considers combinations of evidence.

Can genuine users show no signals?

Yes. Script blockers, accessibility tools, quick reading, and privacy tools can reduce visible signals.

What should I export before deleting data?

Export important evidence, reports, complaint records, and configuration details.

Does renaming the plugin folder delete data?

Normally, no. It temporarily prevents WordPress from loading the plugin. Server-level block rules may remain active until removed.

Why might my IP change?

Dynamic broadband, mobile carriers, VPNs, office networks, and router reconnects can change the public address.

Should I block an entire ASN?

Only with strong evidence and careful review. Large ASNs may contain many legitimate users.

Can BotExorcist prove that a click was charged?

It can record paid-arrival context where detectable, but advertising billing records remain the provider's source for billing confirmation.

Can BotExorcist work without external services?

The core local firewall should continue operating. Optional intelligence or AI-assisted services may require external access.

What happens if an external service is unavailable?

The related optional feature may be unavailable, but the free core firewall should continue operating.

Does BotExorcist send executable code from its server?

BotExorcist Lite must not download or execute externally supplied plugin code.

Is BotExorcist a legal service?

No. It provides evidence and operational tools, not legal advice.

---

51. Glossary

ASN
Autonomous system number. Identifies a network or routing organisation.

Blacklist
A list of addresses or networks explicitly denied.

Bot
Automated software that performs requests or interactions.

CIDR
A notation used to describe an IP network range.

Click signal
A recorded click or activation event.

Crawler
Automated software that browses or indexes content.

Dwell
Time-related engagement on a page or session.

Enforcement
An active response such as blocking or returning a restricted HTTP status.

Evidence Ledger
A structured record of security, behavioural, network, and PPC events.

Governor
The central BotExorcist authority responsible for enforcement decisions.

Honeypot
A deceptive element or route intended to attract automated activity.

HTTP 403
Access is forbidden.

HTTP 404
The requested resource was not found.

HTTP 410
The requested resource is intentionally gone.

Human signal
Interaction that may indicate genuine human engagement.

IP address
A network address associated with a request.

Observation Mode
A mode that records activity without immediate enforcement.

PPC
Pay-per-click advertising.

Query string
The part of a URL containing parameters after ?.

Referrer
The reported source page or site.

Risk score
An accumulated behavioural assessment.

Route
A requested path or endpoint.

Safe lane
A narrowly defined exception for a legitimate WordPress process.

Whitelist
A list of addresses, networks, or requests explicitly trusted.

---

52. Appendix A: event taxonomy

The exact event names depend on version.

A documentation entry for each event should include:

- event name;
- plain-language meaning;
- likely causes;
- risk relevance;
- score effect;
- enforcement relevance;
- administrator exclusion;
- troubleshooting guidance.

Suggested event categories follow.

52.1 Human events

- click;
- scroll;
- navigation;
- page-two interaction;
- meaningful dwell;
- form interaction.

52.2 Suspicious events

- honeypot;
- unknown route;
- parameter probe;
- repeated 404;
- repeated 410;
- suspicious redirect;
- no-signal arrival;
- rapid request sequence.

52.3 Enforcement events

- observed;
- allowed;
- blocked;
- whitelisted;
- blacklisted;
- 403 issued;
- 410 issued;
- rule written;
- rule removed.

52.4 Administrative events

- manual whitelist;
- manual block;
- override;
- trusted login;
- configuration change;
- reset;
- export.

---

53. Appendix B: configuration reference

Each setting should eventually be documented with:

Field	Description
Setting name	The exact label shown in the plugin
Default	Default value
Accepted values	Available options
Purpose	What the setting controls
Recommended use	Normal recommendation
Risk	Possible downside
Edition	Lite or Pro
Related documentation	Anchor link

Core configuration groups should include:

- mode;
- administrator trust;
- IP detection;
- whitelist;
- blacklist;
- scoring;
- human signals;
- routes;
- query intelligence;
- retention;
- PPC;
- exports;
- login shield;
- external services;
- credits;
- licensing;
- diagnostics.

---

54. Appendix C: data-field reference

Each stored field should eventually document:

- name;
- data type;
- source;
- purpose;
- availability;
- limitations;
- retention;
- privacy relevance.

Data groups include:

- activity log;
- IP catalogue;
- ASN intelligence;
- PPC records;
- evidence ledger;
- query intelligence;
- login trust;
- complaint history;
- settings;
- diagnostics.

---

55. Appendix D: support and security reporting

55.1 Support request

Include:

- BotExorcist version;
- WordPress version;
- PHP version;
- hosting environment;
- proxy or CDN;
- relevant timestamp;
- affected route;
- screenshots;
- diagnostics export;
- steps to reproduce.

Remove or redact sensitive data before sharing.

55.2 Security report

A security report should include:

- affected version;
- vulnerability description;
- reproduction steps;
- expected impact;
- proof of concept where safe;
- reporter contact information.

Do not publish an unpatched vulnerability publicly before responsible disclosure has been completed.

55.3 False-positive report

Include:

- IP or network;
- timestamp;
- route;
- event;
- expected behaviour;
- actual behaviour;
- whether the user was authenticated;
- whether a proxy was present;
- whether Observation Mode resolves the issue.

---

56. Audience and recommended reading paths

This manual serves several audiences. Each audience needs a different route through the documentation.

56.1 Website owners

Website owners usually need to understand:

- what BotExorcist protects;
- what remains free;
- what credits are used for;
- how to install safely;
- how to avoid administrator lockout;
- how to read the dashboard;
- how to investigate a block;
- how to recover access;
- how to obtain support.

Recommended reading:

1. BotExorcist Lite
2. Credits and optional services
3. Observation Mode
4. Installation
5. Initial configuration
6. Understanding the dashboard
7. Emergency recovery

56.2 Agencies and PPC managers

Agencies and advertising specialists usually need to understand:

- how paid arrivals are recognised;
- how engagement is measured;
- how repeat IP and ASN patterns are interpreted;
- what an evidence pack contains;
- how to communicate uncertainty;
- how to separate low-quality traffic from proven fraud;
- how to report findings to clients.

Recommended reading:

1. PPC Integrity
2. Advertising evidence and interpretation
3. Evidence Ledger
4. Reports and exports
5. Complaint and dispute support
6. PPC investigation playbook

56.3 WordPress administrators and developers

Technical administrators usually need to understand:

- request flow;
- the Governor;
- real visitor IP detection;
- safe lanes;
- route and query handling;
- server-level rules;
- performance;
- data storage;
- diagnostics;
- emergency disablement.

Recommended reading:

1. How the firewall works
2. The Governor
3. Routes, paths, and probing
4. Query and parameter intelligence
5. Hosting, CDN, proxy, and caching environments
6. Compatibility and acceptance-test matrix
7. Diagnostics and support bundle

56.4 Support personnel

Support personnel should begin with:

- the installed version;
- the current operating mode;
- the detected administrator IP;
- the reported visitor IP;
- recent relevant activity;
- exact HTTP response;
- server type;
- proxy or CDN;
- reproduction steps;
- whether the problem disappears in Observation Mode.

Support should avoid guessing from a screenshot alone when a diagnostics bundle or activity record is available.

56.5 Security reviewers and marketplace reviewers

A reviewer evaluating BotExorcist Lite should be able to confirm:

- Lite is independently functional;
- the core firewall is not disabled at zero credits;
- Lite does not install Pro;
- Lite does not update Pro;
- the upgrade link opens an ordinary external webpage;
- premium code is downloaded and manually uploaded by the user;
- external services are documented;
- remote services return data, not executable plugin code;
- the product's free purpose is prominent and genuine.

---

57. Product architecture and trust boundaries

BotExorcist is easier to understand when divided into clear trust boundaries.

57.1 Local WordPress boundary

The local WordPress installation is responsible for functions such as:

- receiving the request;
- identifying the route;
- evaluating local settings;
- checking whitelists and blacklists;
- recognising authenticated administrators;
- recording local activity;
- calculating local risk;
- applying local enforcement;
- writing approved server rules;
- presenting the dashboard.

This boundary remains relevant even when no external service is available.

57.2 Browser interaction boundary

Browser-side code may record:

- scroll activity;
- clicks;
- page visibility;
- dwell;
- internal navigation;
- human-interaction combinations.

Browser signals can fail or be absent because of:

- JavaScript blocking;
- privacy extensions;
- browser errors;
- consent restrictions;
- caching;
- accessibility tools;
- network interruption.

For this reason, browser signals are evidence rather than absolute identity proof.

57.3 Server and hosting boundary

The web server may enforce rules before WordPress loads.

This boundary can include:

- Apache `.htaccess`;
- Nginx configuration;
- LiteSpeed rules;
- hosting-provider firewall rules;

- reverse-proxy controls;
- CDN firewall controls.

Server-level enforcement can remain active even when the WordPress plugin folder is renamed. Emergency instructions therefore address both the plugin and the server rule.

57.4 Optional external-service boundary

An optional external service may provide:

- IP enrichment;
- ASN information;
- AI-assisted analysis;
- premium report generation;
- licence or entitlement validation;
- other declared processing.

The external-service boundary must not be confused with plugin distribution.

BotExorcist Lite must not use that boundary to fetch executable premium code.

57.5 Human decision boundary

The administrator remains responsible for high-impact decisions such as:

- blocking a large CIDR;
- blocking an ASN;
- accusing a party of fraud;
- submitting a formal dispute;
- sharing evidence externally;
- changing retention;
- disabling administrator safeguards.

BotExorcist assists human judgement. It does not replace it.

57.6 Trust precedence

A safe general precedence is:

1. Emergency recovery and explicit administrator protection
2. Valid whitelist or trusted session
3. Narrow system safe lane
4. Observation rules
5. Accumulated behavioural and network evidence
6. Governor enforcement decision
7. Logging and audit trail

A lower-precedence rule should not silently defeat a higher-precedence safety rule.

57.7 Failure behaviour

The product should fail safely.

Examples:

- If optional AI analysis fails, local firewall protection continues.
- If IP enrichment fails, the request is not automatically treated as malicious.
- If JavaScript signals are missing, the absence is recorded but does not alone prove automation.

- If the database cannot record an optional dashboard event, the administrator should receive a diagnostic warning.
- If the visitor IP cannot be determined confidently behind a proxy, strict network-wide enforcement should be avoided.
- If a migration is incomplete, the plugin should avoid corrupting evidence or repeatedly applying destructive changes.

---

58. Detailed behavioural scoring reference

The behavioural score converts observations into a structured risk assessment.

The precise values are version-dependent. The following table documents the current reference profile and must be verified against the release being published.

58.1 Reference risk additions

Event	Reference adjustment	Interpretation
One-second or no-signal visit	+3	Weak evidence of drive-by or automated behaviour
Honeypot interaction	+10	Strong evidence that hidden or deceptive content was accessed
Suspicious parameter probe	+5	Query behaviour consistent with probing or pollution
Suspicious 404 event	+5	Request for a nonexistent route with risk context
Suspicious redirect event	+3	Redirect behaviour that adds to an existing pattern

A score addition should not be interpreted without the event context.

For example, a 404 generated by a broken internal link is not equivalent to a sequence of exploit-path requests.

58.2 Reference human-signal reductions

Human event	Reference adjustment	Interpretation
Click	-3	The visitor interacted with an element
Scroll	-1	The visitor explored page content
Scroll and click combination	-6	Stronger combined engagement
Internal navigation	-3	The visitor moved through the site
Page-two scroll and click	-10	Strong evidence of sustained human interest

Human reductions are not permanent guarantees of trust.

A sophisticated automated browser may simulate interaction. Human signals should be combined with route, timing, network, and historical evidence.

58.3 Score before and score after

Every scored event should preserve:

- score before the event;
- event adjustment;

- score after the event;
- reason;
- source module;
- whether the event was automatic or manual;
- whether enforcement changed.

Example:

```
Score before: 8
Event: Scroll and click
Adjustment: -6
Score after: 2
Outcome: Continue observation
```

58.4 Score floors and ceilings

The implementation should define whether:

- scores may become negative;
- scores are capped;
- old events decay;
- trusted events reset a score;
- scores are session-based or persistent;
- score thresholds differ by context.

The documentation must match the actual code.

58.5 Session score versus historical risk

A visitor can have more than one relevant assessment:

- current-session behaviour;
- historical IP record;
- historical network pattern;
- manual classification;
- current request severity.

A good decision does not blindly combine all history into one permanent label.

58.6 Redirect allowance

Where a redirect allowance is configured, ordinary navigation redirects should not be treated as suspicious until the allowance is exceeded or combined with other risk.

The current reference allowance is five redirects. This must be confirmed against the published version.

58.7 Example scoring journey: genuine visitor

1. Paid visitor arrives and initially shows no signal: `+3`.
2. Visitor scrolls: `-1`.
3. Visitor clicks a service card: `-3`.
4. Visitor opens a second page and interacts: up to `-10` for the qualifying page-two combination.
5. The score is redeemed and enforcement is not justified.

58.8 Example scoring journey: automated probe

1. Unknown exploit-like route: +5.
2. Suspicious parameter: +5.
3. Honeypot: +10.
4. No human signals: +3.
5. Repeated requests continue.

The accumulated pattern is substantially stronger than any one event.

58.9 Example scoring journey: ambiguous monitoring service

1. Data-centre IP arrives.
2. No scroll or click is recorded.
3. The request accesses a public page only.
4. Requests occur at a predictable low frequency.
5. User agent and reverse lookup are consistent with a legitimate monitor.

The correct response may be observation or a narrow whitelist, not an automatic ASN block.

58.10 Manual overrides

An administrator should be able to record:

- trusted;
- blocked;
- do not block;
- investigate;
- temporary exception;
- permanent exception.

Manual overrides should preserve the original evidence rather than deleting history.

---

59. Governor decision workflow

The Governor should follow an explainable decision process.

59.1 Step one: determine request identity and context

Gather:

- detected IP;
- proxy confidence;
- authenticated status;
- administrator status;
- route;
- method;
- query;
- referrer;
- campaign context;
- current mode.

If the source IP is unreliable, network-based enforcement should be conservative.

59.2 Step two: apply absolute safety checks

Check:

- emergency bypass;
- administrator IP;
- authenticated administrator;
- explicit whitelist;
- valid trusted session;
- approved safe lane.

When an absolute safety rule applies, the request should be allowed unless the safety rule itself has been invalidated by a deliberate administrator action.

59.3 Step three: classify the route

Determine whether the request is:

- public content;
- asset;
- admin request;
- login request;
- AJAX;
- REST;
- cron;
- builder preview;
- builder save;
- checkout;
- webhook;
- unknown route;
- known probe target.

59.4 Step four: collect relevant evidence

Evidence may include:

- current score;
- route severity;
- query severity;
- IP history;
- ASN history;
- human signals;
- no-signal state;
- request frequency;
- PPC arrival state;
- manual classification;
- previous enforcement.

59.5 Step five: respect operating mode

Observation Mode

- record;
- score;
- classify;
- do not enforce ordinary first-time uncertainty;
- preserve severe-event handling only where explicitly designed.

Enforcement Mode

- apply configured thresholds;
- respect all trusted exclusions;
- select the narrowest appropriate action;
- record the explanation.

59.6 Step six: select an action

Possible actions:

- allow;
- allow and observe;
- allow but flag;
- return 403;
- return 410;
- write a temporary block;
- write a persistent block;
- request manual review.

59.7 Step seven: explain the decision

Every high-impact decision should answer:

- What happened?
- Which evidence mattered?
- Which safety checks were evaluated?
- What score changed?
- Which rule was applied?
- Was the action temporary or persistent?
- How can the administrator reverse it?

59.8 Step eight: update related records

After the decision, update only the relevant systems:

- activity log;
- IP record;
- ASN record;
- PPC record;
- Evidence Ledger;
- query history;
- complaint history;
- dashboard counters.

A dashboard reset should not unintentionally erase unrelated forensic records.

59.9 Decision examples

Whitelisted administrator triggers a suspicious parameter rule

Expected result:

- whitelist wins;
- request is allowed;
- event may be recorded for diagnostics;
- no block is applied.

New paid visitor shows no interaction

Expected result:

- observe;
- record low-signal state;
- do not immediately declare fraud;
- wait for more evidence.

Known hostile source requests several exploit routes

Expected result:

- accumulate route evidence;
- consider immediate 403 or server block;
- preserve event sequence;
- avoid exposing sensitive implementation detail.

Page builder performs a valid AJAX save

Expected result:

- authenticated administrator and safe lane are recognised;
- request is allowed;
- suspicious route logic does not override the editor lane.

---

60. PPC Integrity metrics dictionary

PPC Integrity dashboard cards need plain-language definitions so users do not overinterpret them.

60.1 Unique paid IPs

The number of distinct IP addresses associated with recognised paid arrivals during the selected period.

Limitations:

- one visitor can use several IPs;
- several visitors can share one IP;
- VPN and mobile traffic can change addresses;
- an address is not a person.

60.2 Bot behaviour

The number or proportion of paid sessions containing events associated with automation or suspicious behaviour.

This may include:

- honeypot activity;
- suspicious routes;
- suspicious parameters;
- repeated no-signal behaviour;
- rapid requests;
- known automation patterns.

It is not automatically equal to confirmed invalid clicks.

60.3 Human signals

The amount of recorded engagement, such as:

- scroll;
- click;
- navigation;
- page depth;
- dwell;
- combined interaction.

The metric should specify whether it counts events, sessions, or unique visitors.

60.4 Honeypot

The number of paid or related arrivals that interacted with a hidden or deceptive element.

Honeypot events are usually stronger evidence than simple non-engagement.

60.5 Scrolls

The number of qualifying scroll events or sessions.

The interface should state:

- minimum scroll threshold;
- whether repeated scrolls are deduplicated;
- whether the value represents people, sessions, or events.

60.6 Clicks

The number of qualifying click events or sessions.

Clicks on administrative or non-meaningful elements should not inflate the metric.

60.7 Drive-by visits

Paid arrivals with a very short or low-signal session.

The exact rule should be documented.

A drive-by visit may be:

- accidental;
- irrelevant;
- automated;
- interrupted;
- caused by poor page performance.

60.8 Charged without interest

A paid arrival that produced no qualifying evidence of genuine interest.

This wording should be presented as an analytical category, not a statement of proven provider wrongdoing.

60.9 Repeat IPs

Addresses associated with more than one relevant paid arrival.

The report should include:

- number of arrivals;
- date span;
- campaigns;
- landing pages;
- signals;
- enforcement;
- ASN.

60.10 Repeat ASNs

Networks appearing repeatedly in paid traffic.

A repeat ASN is a lead for investigation, not a reason to block a large network automatically.

60.11 Top ASN

The ASN with the greatest relevant activity for the selected period.

The dashboard should show both volume and context.

60.12 Top city

The most frequently resolved approximate city.

Because IP geolocation is approximate, this card should not be treated as exact physical-location evidence.

60.13 Provider-specific totals

Google, Microsoft, and Meta data should remain distinguishable because:

- parameters differ;
- attribution differs;
- campaign structures differ;
- dispute processes differ.

60.14 Monthly repeat analysis

A monthly view helps distinguish:

- one-time anomalies;
- persistent network patterns;
- campaign-specific changes;
- seasonal traffic;
- new bot infrastructure.

60.15 Date-range integrity

Every dashboard and export should clearly show:

- start date;
- end date;
- time zone;
- whether the current day is partial;
- whether data retention truncates the range.

---

61. PPC investigation playbook

A PPC investigation should follow a repeatable process rather than beginning with an accusation.

61.1 Define the question

Examples:

- Why did leads decline while spend remained stable?
- Is one campaign attracting unusually low engagement?
- Are repeated networks consuming paid traffic?
- Did a specific period contain bot-like activity?
- Is one landing page associated with drive-by visits?
- Is the apparent problem attribution, page performance, or traffic quality?

61.2 Set the evidence window

Choose a date range that is:

- long enough to reveal patterns;
- short enough to preserve campaign relevance;
- aligned with billing and campaign reports;
- within the local retention period.

Record the site time zone and advertising-account time zone.

61.3 Preserve the baseline

Before filtering suspicious activity, record:

- total provider clicks;
- total recognised paid arrivals;
- conversions;
- spend;
- campaign changes;
- landing-page changes;
- website outages;
- tracking changes;
- consent-banner changes.

61.4 Segment the data

Useful segments include:

- provider;
- campaign;
- keyword or search term where available;
- landing page;
- device;
- country;
- city;
- IP;

- ASN;
- hour of day;
- day of week;
- engagement class;
- human-signal class.

61.5 Identify clusters

Look for combinations such as:

- repeated IP plus repeated no-signal visits;
- repeated ASN plus many changing IPs;
- one landing page plus abnormal bounce behaviour;
- one campaign plus suspicious parameter probes;
- paid arrivals followed by exploit routes;
- high click count with unusually low recognised arrivals;
- one geography inconsistent with targeting.

61.6 Compare suspicious and normal traffic

A useful investigation compares:

- dwell;
- scroll;
- clicks;
- page depth;
- routes;
- conversion events;
- network types;
- timing.

This prevents a conclusion based only on a suspicious subset.

61.7 Grade the evidence

A practical internal grading model:

Grade A: strong multi-signal evidence

Examples:

- honeypot;
- exploit probing;
- repeated paid arrivals;
- no human signals;
- repeated network pattern.

Grade B: persuasive pattern

Examples:

- repeated IP or ASN;
- repeated low-signal sessions;
- consistent timing;
- no conversion;
- limited route activity.

Grade C: anomaly requiring review

Examples:

- single drive-by;
- one unusual city;
- one data-centre address;
- one no-scroll session.

Grade D: insufficient evidence

Examples:

- one short visit;
- missing JavaScript;
- one shared IP;
- geolocation mismatch alone.

61.8 Investigate alternative explanations

Before escalating, ask:

- Was the landing page slow?
- Was the offer relevant?
- Did tracking fail?
- Did the consent manager block signals?
- Was the campaign broad match?
- Was the visitor using a privacy browser?
- Did the site redirect incorrectly?
- Was a monitoring service involved?
- Did a page-builder or cache issue remove scripts?

61.9 Prepare the report

A responsible report should contain:

1. Executive summary
2. Scope and date range
3. Data sources
4. Method
5. Key findings
6. Evidence tables
7. Alternative explanations
8. Limitations
9. Recommended actions
10. Supporting exports

61.10 Recommended actions

Actions may include:

- exclude poor placements;
- tighten targeting;
- revise keywords;
- block a narrow confirmed range;
- adjust landing pages;
- repair tracking;
- preserve evidence;

- contact a provider;
- continue observation;
- avoid action until more evidence exists.

61.11 Provider dispute package

A dispute package may contain:

- account and campaign identifiers;
- billing period;
- provider click report;
- BotExorcist recognised-arrival report;
- repeat IP and ASN tables;
- representative event records;
- methodology;
- explicit limitations;
- requested review.

The language should remain factual.

---

62. Evidence quality and chain of custody

Evidence is most useful when its origin, handling, and limitations are clear.

62.1 Source integrity

For each report, record:

- website;
- BotExorcist version;
- export time;
- site time zone;
- selected date range;
- active filters;
- relevant configuration;
- data-retention state.

62.2 Preserve originals

Keep an unmodified original export.

Create a separate working copy for:

- sorting;
- highlighting;
- commentary;
- calculations;
- client presentation.

62.3 Hashes and file integrity

For high-value investigations, calculate a cryptographic hash of the original export.

Example record:

```
File: evidence-ledger-2026-07-01-to-2026-07-15.csv
SHA-256: [recorded hash]
Exported: 2026-07-16 14:25 SAST
BotExorcist version: [version]
```

A hash does not prove the underlying event was true. It helps show that the exported file has not changed since hashing.

62.4 Time integrity

Confirm:

- WordPress time zone;
- server time zone;
- database time handling;
- advertising-account time zone;
- daylight-saving differences where applicable.

Do not compare timestamps without normalising them.

62.5 Evidence completeness

A record may be incomplete because:

- a referrer was withheld;
- campaign parameters were removed;
- JavaScript did not run;
- an external lookup failed;
- the IP was hidden behind a proxy;
- retention removed older events;
- a cache served a page without fresh instrumentation.

Reports should identify missing fields rather than silently treating them as negative evidence.

62.6 Manual notes

Manual notes should preserve:

- author;
- date;
- reason;
- source;
- whether the note is opinion or verified fact.

62.7 Screenshots

Screenshots are useful for presentation but weaker than exported structured data when:

- the date range is not visible;
- filters are hidden;
- totals cannot be reproduced;
- the image is cropped;
- metadata is unavailable.

Use screenshots as supporting material, not the only record.

62.8 Data minimisation when sharing

Before sending evidence externally:

- redact administrator IPs where unnecessary;
- remove customer personal information;
- remove tokens and secrets;
- remove unrelated request parameters;
- limit the export to the relevant period;
- password-protect sensitive archives;
- use a secure delivery method.

62.9 Interpretation record

Keep a distinction between:

- raw event;
- system classification;
- administrator interpretation;
- final conclusion.

Example:

```
Raw event: Paid arrival, no scroll, 1.2-second dwell
System classification: Drive-by / low signal
Administrator interpretation: Suspicious when combined with 18 similar arrivals
Conclusion: Request provider review; not represented as proven fraud
```

---

63. Compatibility and acceptance-test matrix

BotExorcist should not be considered production-ready on a website until the website's critical workflows have been tested.

The goal is not merely to confirm that the homepage opens. The goal is to prove that legitimate requests can complete while suspicious requests remain observable and enforceable.

63.1 Testing stages

Stage one: baseline without active enforcement

Record the normal behaviour of:

- public pages;
- administrator login;
- forms;
- checkout;
- page builders;
- REST requests;
- AJAX;
- scheduled tasks;
- webhooks.

Stage two: Observation Mode

Enable BotExorcist in Observation Mode and repeat the tests.

Review:

- logged events;
- score changes;
- route classifications;
- IP detection;
- safe-lane recognition;
- false-positive indicators.

Stage three: controlled enforcement

Enable enforcement in a staging environment or low-risk period.

Test one critical workflow at a time.

Stage four: production monitoring

After deployment:

- monitor new blocks;
- review administrator activity;
- review forms and orders;
- confirm PPC tracking;
- preserve a rollback route.

63.2 Core acceptance matrix

Area	Test	Expected result
Homepage	Open as logged-out visitor	Page loads and normal assets are allowed
Internal page	Navigate through menu	Human navigation is recorded
Administrator login	Log in successfully	Administrator becomes trusted
Failed login	Enter incorrect password	Event is recorded without harming legitimate recovery
Password reset	Request reset email	Request completes
Logout	Log out	Session ends normally
Admin dashboard	Open main screens	No 403 or 410 responses
Plugin activation	Activate and deactivate	No fatal errors
Permalinks	Save permalink settings	`.htaccess` remains valid
Media	Upload an image	Upload and AJAX requests complete
REST API	Open a known endpoint	Valid endpoint remains accessible
AJAX	Trigger a normal AJAX action	Request passes safe-lane checks
WP-Cron	Run scheduled task	Cron request is not blocked
Search	Submit a site search	Legitimate query parameters are allowed
404 page	Request a random missing page	Correct 404 handling occurs
Junk URL	Request a confirmed polluted URL	Configured 410 behaviour occurs
Whitelist	Add test IP	Whitelist overrides ordinary enforcement
Manual block	Block test address	Correct 403 or server block occurs
Unblock	Remove test block	Access is restored fully

63.3 Page-builder matrix

Test each installed builder in the modes actually used by the site.

Builder workflow	Test
Open visual editor	Editor loads completely
Open preview	Preview route is allowed
Save draft	Save request completes
Publish update	Publish request completes
Regenerate CSS	Dynamic files are written
Open template library	External and REST calls work
Edit reusable template	AJAX and REST requests work
Exit editor	Normal frontend page loads

When a request is exempted, the rule should be narrow enough to cover the builder workflow without exempting unrelated public traffic.

63.4 WooCommerce matrix

Test:

- product page;
- variation selection;
- add to cart;
- cart update;
- coupon;
- checkout;

- account login;
- account registration;
- payment redirect;
- payment return;
- webhook;
- order email;
- subscription renewal if present;
- refund callback if present.

A payment callback may not show human interaction. It should be trusted by its validated route, signature, provider range, or integration logic—not by browser signals.

63.5 Form matrix

For every important form, test:

- anonymous submission;
- logged-in submission;
- file upload;
- CAPTCHA;
- AJAX submission;
- confirmation redirect;
- email notification;
- CRM webhook;
- validation error;
- repeated legitimate submission.

63.6 Caching and optimisation matrix

Test with the actual production combination of:

- page cache;
- object cache;
- minification;
- deferred JavaScript;
- CDN;
- image optimisation;
- consent manager;
- reverse proxy.

Verify that BotExorcist's browser-side signals are not removed, delayed beyond usefulness, or served with stale configuration.

63.7 Proxy and CDN matrix

Verify:

- direct origin requests;
- proxied public requests;
- administrator requests;
- IPv4;
- IPv6 where available;
- Cloudflare or CDN header;
- spoofed forwarded header;
- correct trusted-proxy range;
- fallback when header is missing.

The test must prove that an arbitrary visitor cannot choose a trusted IP merely by sending a forged header.

63.8 Security regression matrix

A release should also confirm that:

- whitelists still override blocks;
- administrator protection still applies;
- Observation Mode does not issue ordinary first-time 403 responses;
- Lite does not fetch executable Pro code;
- external-service failure does not disable the local firewall;
- server rules can be removed;
- emergency recovery still works;
- reset actions do not erase unrelated data unexpectedly;
- query overrides persist;
- complaint cooldowns persist;
- evidence exports contain the selected period only.

63.9 Acceptance-test record

Maintain a record such as:

```
Site:
Environment:
BotExorcist version:
WordPress version:
PHP version:
Server:
CDN/proxy:
Cache:
Test date:
Tester:

Workflow:
Expected:
Actual:
Pass/fail:
Relevant event ID:
Notes:
```

---

64. Emergency recovery by access method

The core emergency procedure is described in section 41. This section provides practical variants for common hosting environments.

64.1 cPanel File Manager

1. Open cPanel.
2. Open **File Manager**.
3. Enable **Show Hidden Files** so `.htaccess` is visible.
4. Open the website's document root, often `public_html`.
5. Download `.htaccess` as a backup.
6. Edit `.htaccess` and remove only the rule blocking your IP.
7. Open `wp-content/plugins/`.
8. Rename the active folder to `.off`.
9. Log in to WordPress.
10. Restore the folder name.
11. Reactivate if required.
12. Whitelist the administrator IP.
13. return to Observation Mode for testing.

64.2 Plesk File Manager

1. Open the relevant domain.
2. Select **Files**.
3. Open the WordPress document root.
4. Back up `.htaccess`.
5. Remove the specific BotExorcist deny rule.
6. Open `wp-content/plugins`.
7. Rename the plugin folder.
8. Restore WordPress access.
9. Return the folder to its original name.
10. correct IP detection and whitelist settings.

64.3 FTP or SFTP

1. Connect using an FTP or SFTP client.
2. Enable display of hidden files.
3. Download `.htaccess`.
4. Edit it locally with a plain-text editor.
5. Remove the relevant BotExorcist block rule.
6. Upload the corrected file.
7. Rename the plugin directory remotely.
8. Log in to WordPress.
9. Restore the directory name.
10. Retest safely.

Use SFTP rather than unencrypted FTP where the host supports it.

64.4 SSH

Example navigation:

```
cd /path/to/wordpress
cp .htaccess .htaccess.botexorcist-backup
cd wp-content/plugins
mv botexorcist botexorcist.off
```

After removing the specific block from `.htaccess`, log in and correct the configuration.

Restore:

```
mv botexorcist.off botexorcist
```

For Lite:

```
mv botexorcist-lite botexorcist-lite.off
mv botexorcist-lite.off botexorcist-lite
```

Do not copy commands blindly when the actual WordPress path or folder name differs.

64.5 WP-CLI deactivation

If SSH and WP-CLI remain available, the plugin may be deactivated through WordPress.

Examples:

```
wp plugin deactivate botexorcist
```

or:

```
wp plugin deactivate botexorcist-lite
```

This does not necessarily remove a server-level `.htaccess` block. The blocking rule may still need to be removed separately.

64.6 Nginx servers

Nginx does not use `.htaccess`.

A block may exist in:

- site configuration;
- included BotExorcist rule file;
- hosting firewall;
- reverse proxy;
- control-panel deny list.

On Nginx:

1. Disable or remove the exact deny rule.
2. Validate the configuration.
3. Reload Nginx.
4. Rename or deactivate the plugin if needed.
5. Correct the source of the block.

Example validation and reload commands vary by host and should be performed by an authorised administrator.

64.7 Cloudflare or external firewall

If the WordPress files are correct but access remains blocked, check:

- Cloudflare WAF events;
- IP Access Rules;
- custom firewall rules;
- rate-limiting rules;
- hosting-provider firewall;
- ModSecurity;
- fail2ban;
- control-panel deny list.

BotExorcist cannot restore access if another security layer continues to deny the address.

64.8 Dynamic IP changed during recovery

If the administrator's public IP changed:

- identify the current public IP;
- remove the old and current block where necessary;
- avoid permanently whitelisting a broad mobile range;
- prefer trusted authenticated-session logic;
- review strict IP lock settings.

64.9 Recovery verification

After recovery, verify all of the following:

- frontend opens;
- `/wp-admin/` opens;
- login succeeds;
- logout succeeds;
- current IP is correct;
- whitelist is present;
- `.htaccess` contains valid syntax;
- no duplicate deny rule remains;
- plugin is in Observation Mode;
- emergency file access remains available.

---

65. Diagnostics and support bundle

A support request should be reproducible without requiring unrestricted access to the user's website.

65.1 Minimum diagnostics

The diagnostics page should report:

- BotExorcist edition;
- BotExorcist version;
- database schema version;
- WordPress version;
- PHP version;
- web server;
- server API;
- operating mode;
- site URL;
- home URL;
- multisite state;
- detected administrator IP;
- real-IP source header;
- proxy or CDN state;
- object cache;
- page cache where detected;
- active theme;
- relevant installed integrations;
- relevant writable paths;
- scheduled-task status.

65.2 Firewall state

Include:

- Governor enabled or disabled;
- Observation or Enforcement Mode;
- whitelist count;
- blacklist count;
- CIDR rule count;
- `.htaccess` management state;
- last server-rule write;
- last rule error;
- current emergency bypass state;
- administrator trust state.

65.3 Database state

Include:

- required table presence;
- table versions;
- approximate row counts;
- last cleanup;

- retention setting;
- migration status;
- database error summary.

Do not include database passwords.

65.4 External-service state

Include only safe diagnostic information:

- service enabled;
- last successful request;
- last error category;
- timeout;
- remaining credits where appropriate;
- entitlement state;
- endpoint hostname.

Do not export API secrets, full licence keys, or authentication tokens.

65.5 Recent-event extract

A support bundle may include a limited recent event window surrounding the reported problem.

Useful fields:

- timestamp;
- route;
- method;
- action;
- event;
- score before;
- score after;
- administrator state;
- safe-lane result;
- IP masked or redacted where appropriate;
- error code.

65.6 Redaction

The export process should redact:

- passwords;
- cookies;
- authorization headers;
- nonces where sensitive;
- API keys;
- payment data;
- private form data;
- full licence keys;
- session tokens.

65.7 Diagnostics export format

A useful bundle may contain:

```
botexorcist-diagnostics/
├── summary.json
├── environment.txt
├── firewall-state.json
├── recent-events.csv
├── table-health.txt
├── scheduled-tasks.txt
└── README.txt
```

The README should explain:

- what the bundle contains;
- what has been redacted;
- creation time;
- BotExorcist version;
- how the user can review it before sharing.

65.8 One-click copy summary

The Help page should provide a short copyable summary for support tickets.

Example:

```
BotExorcist Pro 2.x
WordPress 6.x
PHP 8.x
Server: LiteSpeed
Mode: Observation
Proxy: Cloudflare
Detected IP source: CF-Connecting-IP
Issue: Elementor save returns 403
First observed: [timestamp]
```

65.9 Health warnings

The plugin should surface warnings for:

- missing database table;
- failed migration;
- unwritable server-rule file;
- incorrect real-IP detection;
- administrator not whitelisted;
- expired cleanup task;
- external-service timeout;
- unusually large log table;
- repeated fatal errors;
- conflicting block rules.

---

66. Documentation publishing and in-plugin integration

The Markdown file is the master content source.

The public website, plugin help, purchase instructions, and PDF manual should be derived from it rather than maintained as unrelated copies.

66.1 Public documentation URL

Recommended canonical location:

```
https://botexorcist.com/documentation/
```

Useful related routes may include:

```
/documentation/getting-started/
/documentation/firewall/
/documentation/ppc-integrity/
/documentation/troubleshooting/
/documentation/emergency-recovery/
/documentation/glossary/
```

The first publication may be one long page with anchor navigation. It can later be split into chapter pages without changing the underlying content source.

66.2 Page structure

A practical desktop layout:

- compact brand header;
- documentation title;
- search;
- left sticky table of contents;
- main reading column;
- right-side “On this page” menu for long chapters;
- previous and next chapter controls;
- print control;
- version indicator.

Mobile layout:

- collapsible contents panel;
- full-width reading column;
- persistent search;
- back-to-top control.

66.3 Documentation search

The search index should include:

- headings;
- body text;
- dashboard labels;

- setting names;
- HTTP codes;
- glossary terms;
- common error messages;
- emergency phrases such as “locked out” and “403”.

Search results should link to the exact section.

66.4 Anchor stability

Published anchors should remain stable.

If a heading is renamed, preserve an alias or redirect so existing in-plugin links do not break.

66.5 Version labels

Each page should state:

- documentation last updated;
- applicable BotExorcist version or range;
- Lite, Pro, or both;
- whether a feature is experimental;
- whether exact settings differ by version.

66.6 Built-in local excerpts

The plugin should embed essential local help for:

- emergency recovery;
- administrator IP;
- Observation Mode;
- whitelist;
- unblocking;
- credits;
- Lite versus Pro;
- safe-mode troubleshooting.

This local content should be packaged with the plugin and not fetched dynamically as executable or privileged content.

66.7 External links

A contextual link should open the exact documentation section.

Example:

```
$docs_url =
 ↪  'https://botexorcist.com/documentation/#41-emergency-recovery-blocked-or-locked-out';
```

The link is ordinary navigation. It is not a remote installer.

66.8 In-plugin help cards

A standard help card should contain:

- title;
- one-paragraph explanation;
- immediate action;

- risk warning where relevant;
- “Read full documentation” link.

Example:

```
Observation Mode

BotExorcist records and scores traffic without applying ordinary active blocks.
Use this mode during installation, migrations, page-builder testing, and proxy changes.

Action: Keep Observation Mode enabled until login, forms, checkout, AJAX, REST, and page
 ↪  builders have been tested.

Read full documentation ->
```

66.9 Content generation pipeline

A sustainable workflow:

```
Master Markdown
      ↓
Validation and product review
      ↓
Website HTML
      ↓
Search index
      ↓
Plugin-local excerpts
      ↓
Printable PDF
      ↓
Support and purchase snippets
```

66.10 Documentation governance

Every feature release should identify:

- affected documentation sections;
- new settings;
- removed settings;
- migration changes;
- new events;
- new dashboard cards;
- new troubleshooting entries;
- changed Lite/Pro availability.

A feature is not complete until its documentation is updated.

---

67. Purchase, onboarding, and transactional instructions

The customer should receive critical safety information before activation, not after a lockout.

67.1 Purchase-page product clarification

Recommended wording:

> ***BotExorcist Lite remains a free firewall.***
> *BotExorcist Pro is a separate premium product with advanced traffic intelligence, evidence, PPC Integrity, reporting, and investigation capabilities. Purchasing Pro does not change the continued operation of the Lite core firewall.*

67.2 Credit clarification

Recommended wording:

> ***Credits do not power the core firewall.***
> *Credits apply only to identified optional services such as enhanced intelligence or AI-assisted processing. If credits reach zero, the local core firewall continues operating.*

67.3 Installation warning

Recommended wording:

> ***Before installation:*** *Confirm that you can access your hosting file manager, FTP, SFTP, or SSH. Begin in Observation Mode and confirm your administrator IP before enabling active enforcement.*

67.4 Emergency recovery card

Recommended compact wording:

> ***Blocked or locked out?***
>
> 1. *Remove your IP from the BotExorcist block in* `.htaccess`*.*
> 2. *Rename* `/wp-content/plugins/botexorcist` *to* `botexorcist.off`*.*
> 3. *Log in to WordPress normally.*
> 4. *Rename the folder back.*
> 5. *Reactivate if required.*
> 6. *Whitelist your IP and verify real-IP detection.*
>
> *For Lite, use* `botexorcist-lite.off`*.*

67.5 Order-confirmation email

Suggested content:

```
Thank you for purchasing BotExorcist Pro.

Your next steps:
```

```
1. Download the BotExorcist Pro ZIP from your account.
2. Back up your WordPress files and database.
3. Confirm access to your hosting file manager, SFTP, or SSH.
4. In WordPress, open Plugins -> Add New -> Upload Plugin.
5. Upload the ZIP and activate it.
6. Begin in Observation Mode.
7. Confirm your administrator IP.
8. Test login, forms, checkout, AJAX, REST, and page builders before enabling enforcement.

Emergency recovery instructions:
[direct documentation link]
```

67.6 Download-page checklist

The account download page should display:

- purchased product;
- current version;
- release date;
- minimum WordPress version;
- minimum PHP version;
- changelog;
- checksum where offered;
- installation guide;
- emergency recovery;
- support link.

67.7 First-run onboarding

Recommended onboarding sequence:

Step 1: Environment check

- server;
- WordPress;
- PHP;
- proxy;
- real-IP confidence;
- writable rule file.

Step 2: Administrator protection

- detected IP;
- whitelist;
- authenticated trust;
- emergency access confirmation.

Step 3: Operating mode

- Observation Mode selected by default;
- explanation of why;
- enforcement unavailable until acknowledgement where appropriate.

Step 4: Compatibility tests

- login;
- forms;
- checkout;
- page builder;
- REST;
- AJAX.

Step 5: Optional services

- explain credits;
- explain data transfer;
- allow opt-in;
- show what remains local.

Step 6: Finish

- show Help page;
- show emergency recovery link;
- show diagnostics export;
- show current protection state.

67.8 Zero-credit notice

A zero-credit notice should not imply that the firewall is inactive.

Preferred wording:

> ***Optional service credits are exhausted.***
> *The core BotExorcist firewall remains active. Optional credit-based analysis is paused until credits are added or the service is disabled.*

Avoid wording such as:

- protection expired;
- firewall disabled;
- site unprotected;
- trial ended;

unless that statement is actually true for a specific optional product and is clearly separated from the free firewall.

67.9 Update instructions

Recommended wording:

> *BotExorcist Lite updates through WordPress.org when available there. BotExorcist Pro is purchased and downloaded separately. Where manual Pro updates apply, download the current ZIP and upload it through WordPress. Lite does not install or update Pro.*

---

68. Release management and change control

BotExorcist is a security product. Releases should prioritise predictable behaviour over speed.

68.1 Release stages

Recommended stages:

1. Development
2. Automated static checks
3. Unit or component tests where available
4. Staging installation
5. Migration test
6. Compatibility matrix
7. Security review
8. Documentation review
9. Release candidate
10. Production release
11. Post-release monitoring

68.2 Version contents

A release record should identify:

- plugin version;
- database schema version;
- release date;
- Lite or Pro;
- minimum WordPress version;
- tested WordPress version;
- minimum PHP version;
- migration requirement;
- changed settings;
- changed score logic;
- changed enforcement logic;
- changed remote-service behaviour;
- changed data fields;
- documentation revision.

68.3 High-risk change categories

Changes require extra testing when they affect:

- administrator trust;
- IP detection;
- `.htaccess`;
- block precedence;
- whitelists;
- Governor thresholds;
- Observation Mode;
- database migration;
- query 410 handling;

- page-builder safe lanes;
- REST or AJAX;
- purchase or update boundaries.

68.4 Scoring-change notice

When scoring changes, document:

- previous value;
- new value;
- reason;
- expected effect;
- whether old scores are recalculated;
- whether thresholds change.

68.5 Database migration checklist

Before release:

- install from a clean database;
- upgrade from the previous stable version;
- upgrade from an older supported version;
- repeat migration safely;
- test interrupted migration;
- test rollback limitations;
- confirm indexes;
- confirm no data loss;
- confirm table-prefix support.

68.6 Server-rule checklist

Test:

- existing WordPress rewrite rules;
- empty BotExorcist section;
- many allow and deny rules;
- IPv4;
- IPv6;
- duplicate rule prevention;
- rule removal;
- malformed file recovery;
- whitelist before blacklist;
- uninstall behaviour.

68.7 WordPress.org Lite review boundary

Before packaging Lite, confirm:

- no Pro installer;
- no Pro updater;
- no remote executable-code delivery;
- no misleading locked storefront screens;
- free firewall purpose is prominent;
- upgrade links are restrained;
- external services are documented;
- text domain matches the plugin slug;

- assets and readme follow directory requirements.

68.8 Changelog format

Use consistent headings:

```
Added
Improved
Fixed
Security
Compatibility
Changed
Deprecated
Removed
Documentation
```

Avoid vague entries such as “miscellaneous fixes.”

68.9 Release rollback package

Retain:

- previous stable ZIP;
- database migration notes;
- rollback instructions;
- known incompatibilities;
- support notice;
- emergency disable instructions.

68.10 Post-release observation

After release, monitor:

- fatal errors;
- migration failures;
- lockout reports;
- page-builder failures;
- checkout failures;
- incorrect IP detection;
- false-positive increases;
- unexpected table growth;
- external-service errors.

---

69. Expanded frequently asked questions

Why does BotExorcist not block every suspicious first visit?

Immediate blocking can create false positives. Observation allows the system to collect more evidence, especially for paid traffic and unfamiliar environments.

Why can a visitor's score go down?

Genuine interaction can redeem an initially suspicious session. Scrolls, clicks, navigation, and page-two engagement can reduce risk.

Can a bot fake human signals?

Yes. Human signals strengthen an assessment but are not absolute proof.

Why is a honeypot considered stronger evidence?

A correctly designed honeypot is not intended for ordinary visitors. Interaction therefore carries more weight than simple inactivity.

Why was a 410 used instead of a 404?

A confirmed junk or permanently removed URL may warrant a 410. Ordinary missing content usually remains a 404.

Can a 410 harm a real page?

Yes, if route recognition is wrong. Real pages, assets, editors, and valid parameters must be exempted and tested.

Does BotExorcist change `.htaccess`?

It may manage server-level allow or deny rules where configured and supported. Back up `.htaccess` and preserve emergency file access.

What if my server uses Nginx?

Nginx does not use `.htaccess`. Server-level rules require an Nginx-compatible method or hosting integration.

Why does disabling the plugin not always unblock me?

A server-level block can remain active after WordPress stops loading the plugin. Remove the relevant server rule as well.

Can I use BotExorcist behind Cloudflare?

Yes, but real-IP detection must trust the correct Cloudflare header only from valid Cloudflare proxy ranges.

Should I whitelist Cloudflare IP ranges as visitors?

No. Proxy ranges should be trusted for forwarding-header interpretation, not treated as the actual human visitor.

Why does the same person have several IPs?

Mobile networks, dynamic broadband, VPNs, IPv4/IPv6 changes, and network switching can change the public address.

Why do many people share one IP?

Offices, schools, hotels, mobile carriers, VPNs, and carrier-grade NAT can place many users behind one public address.

Can I block a whole mobile ASN?

That is usually high risk because the ASN may contain many legitimate customers. Use narrow evidence-based action.

Can BotExorcist identify competitor clicks?

It can identify patterns associated with IPs and networks, but it cannot reliably prove the natural person or employer behind every request.

Can BotExorcist see a Google Ads click ID?

Where the parameter reaches the landing page and is retained, BotExorcist may record relevant campaign identifiers. Redirects, consent tools, or sanitisation may remove them.

Why do provider clicks and recognised arrivals differ?

Possible reasons include:

- blocked requests;
- redirects;
- tracking loss;
- consent restrictions;
- duplicate clicks;
- prefetching;
- page-load failure;
- different time zones;
- provider filtering;
- attribution differences.

Can I send a BotExorcist report directly to Google?

You can use it as supporting evidence. Review the report, include provider records, state limitations, and avoid unsupported conclusions.

Does “charged without interest” mean the click was invalid?

No. It means a paid arrival showed no qualifying interest under the configured measurement.

Can a visitor convert without recorded scroll or click?

Yes. A form may be prefilled, a phone number may be copied outside tracked interaction, JavaScript may be blocked, or another conversion path may exist.

Does Observation Mode mean there is no protection?

Observation Mode limits ordinary active enforcement so the system can learn safely. Other hosting and WordPress security layers may still operate.

Can severe events still be handled in Observation Mode?

That depends on the implemented version and configuration. The manual must describe any exceptions explicitly.

How long should I remain in Observation Mode?

Remain there until critical workflows have been tested and enough normal traffic has been observed to understand the environment.

What is the safest way to enable enforcement?

Enable it gradually, during a monitored period, with file access available and the administrator IP confirmed.

Should I delete all old evidence?

Not automatically. Consider dispute windows, privacy obligations, server capacity, and investigation value.

Does resetting the dashboard delete the Evidence Ledger?

Not necessarily. Each reset must state which data categories it affects.

Why is the map city wrong?

IP geolocation often reflects a provider hub or registration location rather than the visitor’s physical city.

Can I report an ASN for abuse?

Yes, where evidence supports it. Use factual timestamps, source IPs, paths, and behaviour. Do not assume the ASN operator is the attacker.

What happens if IP enrichment is unavailable?

Network details may remain incomplete, but local request evaluation and core firewall functionality should continue.

Does the licence server control the Lite firewall?

No. Lite core protection should not depend on a paid licence.

Does a Pro licence expiry uninstall the plugin?

No. Commercial consequences should be clearly documented. The product should not silently delete itself.

Can I use BotExorcist on staging?

Yes, subject to the applicable entitlement. Staging is strongly recommended for updates and compatibility testing.

Can I run BotExorcist with another firewall plugin?

Possibly, but overlapping enforcement can complicate diagnosis. Test rule precedence, login, REST, AJAX, and server writes.

Which firewall blocked a request?

Review:

- BotExorcist log;
- hosting firewall;
- Cloudflare event log;
- ModSecurity;
- other plugin logs;
- server access and error logs.

What should I send support first?

Send the version, environment summary, exact time, affected route, expected result, actual result, and a redacted diagnostics bundle.

Should support receive my full database?

Not as a first step. Provide the smallest redacted dataset required to reproduce the issue.

Can I edit BotExorcist database tables manually?

Direct edits can break relationships and evidence integrity. Use supported controls or a documented recovery process.

Can I remove one accidental block without resetting everything?

Yes. Remove the specific IP or range and preserve unrelated evidence where the interface supports it.

Why did a safe administrator action appear in the log?

Trusted traffic may still be logged for diagnostics while being exempt from enforcement.

Does a whitelist hide all evidence?

It should bypass enforcement, but the product may still record limited diagnostic activity depending on configuration.

What happens during a server outage?

BotExorcist cannot evaluate requests that never reach WordPress. Hosting, CDN, and server monitoring remain important.

Is BotExorcist a replacement for secure coding?

No. Plugins, themes, passwords, permissions, backups, and patching remain essential.

---

70. Operational navigation map

This part of the manual explains how an administrator performs common tasks inside BotExorcist.

The exact wording of a button may differ slightly between Lite, Pro, and later releases. Before publication, every menu name and control label must be checked against the current build.

The current functional areas are expected to include some or all of the following:

Area	Main purpose
Dashboard	High-level traffic, firewall, network, and PPC summary
Activity Log	Chronological request, event, score, and action history
IP Catalogue	Per-IP intelligence, history, classification, and network details
Access Control	Whitelists, blacklists, CIDR ranges, and manual rules
Query Intelligence	Query-string and parameter decisions
PPC Integrity	Paid-arrival behaviour and provider-specific evidence
Evidence Ledger	Structured evidence records and exports
Login Shield	Login trust, IP or ASN protection, and administrator access controls
Settings	Operating mode, retention, service, display, and system configuration
Help & Documentation	Local help, emergency recovery, diagnostics, and online manual links

70.1 Before changing data

Before deleting, blocking, resetting, or clearing a range:

1. Confirm which screen you are on.
2. Confirm the selected date range.
3. Confirm whether the action affects:
    - display only;
    - stored logs;
    - IP history;
    - PPC evidence;
    - active block rules;
    - `.htaccess`;
    - all of the above.
4. Export evidence if it may be needed later.
5. Confirm that the administrator IP is not part of the selected range.
6. Confirm that the range does not contain:
    - Cloudflare;
    - payment providers;
    - a mobile carrier;
    - a hosting provider;
    - an office network;
    - legitimate crawlers;
    - a client network.
7. Use the narrowest possible range.

70.2 The four different meanings of “remove”

Users often say “remove this IP,” but that can mean four different things.

Remove from the current view

The record is hidden by a filter or date-range change. Nothing is deleted.

Delete historical records

Activity or catalogue records are deleted. This does not necessarily remove an active block.

Unblock

The active deny rule is removed. Historical evidence may remain.

Whitelist

The source is explicitly trusted. Historical evidence remains, and the whitelist should override ordinary enforcement.

The interface and documentation must make these distinctions obvious.

70.3 Where to start

Use this decision guide:

- **A visitor is currently blocked:** Go to Access Control or the IP Catalogue and remove the active block.
- **A CIDR range is blocked:** Remove the CIDR rule and confirm the server rule is removed.
- **You want to delete old activity:** Use Activity Log filters and bulk deletion.
- **You want to remove an IP from the Catalogue:** Use IP Catalogue search and record deletion.
- **You want to remove paid-traffic evidence:** Use PPC Integrity or Evidence Ledger date and source filters.
- **A query parameter is blocked:** Use Query Intelligence and change the verdict or remove the override.
- **You are locked out:** Follow the emergency-recovery procedure before using the dashboard.
- **You need to remove everything related to a source:** Use the multi-area cleanup workflow in section 75.

---

71. Activity Log operations

The Activity Log is the chronological operational record.

It may include:

- timestamp;
- IP address;
- route;
- query;
- event;
- score before and after;
- human signals;
- action;
- block state;
- source module;
- administrator state;
- PPC context.

71.1 Open the Activity Log

In WordPress:

1. Open **BotExorcist**.
2. Select **Activity Log**.
3. Confirm the date range before interpreting totals.
4. Clear old filters if the expected records are not visible.

71.2 Find one IP address

1. Open **Activity Log**.
2. Use the search or IP filter.
3. Enter the complete address.
4. Apply the filter.
5. Review:
    - first event;
    - latest event;
    - route sequence;
    - score changes;
    - human signals;
    - current block state.

Searching the log does not change the source's access.

71.3 Find an IPv4 range

Where the screen supports CIDR filtering:

1. Open the range or CIDR filter.
2. Enter a value such as:

```
192.0.2.0/24
```

3. Apply the filter.
4. Review the matched record count.
5. Export before deletion where the records may be important.
6. Select the required records.
7. Choose the bulk action.

Where CIDR filtering is not available:

1. Search the shared prefix, for example:

```
192.0.2.
```

2. Review every result carefully.
3. Do not assume a text-prefix match is equivalent to a mathematically exact CIDR match.
4. Process each result or page of results.
5. Prefer adding proper CIDR filtering to the interface rather than relying on prefix searches long term.

71.4 Find an IPv6 range

IPv6 text prefixes are easy to misinterpret because addresses may be compressed.

Use a proper CIDR-aware filter where available.

Example:

```
2001:db8::/32
```

Do not use a simple text-prefix deletion as a substitute for IPv6 range calculation unless the interface explicitly normalises addresses.

71.5 Delete selected log records

1. Filter the log.
2. Confirm the selection count.
3. Select individual rows or **Select all on this page**.
4. Where available, choose **Select all matching records** to include results beyond the current page.
5. Choose **Delete records**.
6. Read the confirmation message.
7. Confirm only if the action scope is correct.

After deletion:

- the historical rows are removed;
- dashboard totals derived from those rows may change;
- an active block may remain;
- IP Catalogue records may remain;
- PPC evidence may remain;
- `.htaccess` rules may remain.

71.6 Delete records for a date range

1. Set **From** and **To** dates.
2. Confirm the site time zone.
3. Apply any IP, event, provider, or route filter.
4. Check the result count.
5. Export if needed.
6. select all matching records.
7. delete.

A date range should use inclusive or exclusive boundaries consistently. The UI should state whether the end date includes the complete day.

71.7 Delete one event type

Examples:

- no-signal;
- 404;
- 410;
- honeypot;
- scroll;
- click;
- administrator;
- PPC arrival.

Procedure:

1. Select the event filter.
2. Choose the event type.
3. Add a date range.
4. Add an IP or range filter if required.
5. Review the results.
6. Delete only the intended records.

Do not delete a whole event category merely to make dashboard numbers look better.

71.8 Preserve evidence while reducing clutter

Instead of deleting:

- shorten the visible date range;
- archive an export;
- use event filters;
- collapse repeated events;
- apply retention rules;
- hide administrator events from ordinary views.

71.9 Clear versus delete

A button labelled **Clear view** should only clear filters.

A button labelled **Delete records** should delete stored data.

A button labelled **Reset Activity Log** may delete a broad data set and therefore needs a precise confirmation message.

---

72. IP Catalogue operations

The IP Catalogue is the persistent intelligence record for individual addresses.

An IP Catalogue record may contain:

- current classification;
- first seen;
- last seen;
- event count;
- block state;
- whitelist state;
- ASN;
- organisation;
- location;
- notes;
- enrichment;
- PPC activity;
- history.

72.1 Open an IP record

1. Open **BotExorcist -> IP Catalogue**.
2. Search for the address.
3. Select the record.
4. Review the complete history before changing classification.

72.2 Classify an IP

Possible classifications may include:

- trusted;
- suspicious;
- blocked;
- do not block;
- monitoring service;
- crawler;
- administrator;
- unknown.

When changing classification:

1. Record a reason.
2. Preserve who made the change.
3. Preserve the date.
4. Do not delete the original evidence.

72.3 Block an IP from the Catalogue

1. Open the IP record.
2. Select **Block IP**.
3. Choose temporary or persistent blocking where available.
4. Confirm whether the block will:

- update BotExorcist's database;
- update `.htaccess`;
- update another server rule;
- apply immediately.

5. Confirm.
6. Test from a separate address if practical.

Never block the active administrator IP.

72.4 Unblock an IP from the Catalogue

1. Open the record.
2. Select **Unblock** or **Remove block**.
3. Confirm the database block is removed.
4. Confirm the generated server rule is removed.
5. Clear server and CDN caches where relevant.
6. Retest access.
7. Keep the historical record unless there is a separate reason to delete it.

72.5 Delete one IP Catalogue record

1. Open the IP record.
2. Export or note important evidence.
3. Choose **Delete record**.
4. Read the confirmation scope.
5. Confirm.

Deleting the Catalogue record should not be assumed to:

- remove Activity Log rows;
- remove PPC rows;
- remove Evidence Ledger rows;
- remove a blacklist entry;
- remove `.htaccess` rules.

The interface should either coordinate these actions or explicitly state what remains.

72.6 Delete several Catalogue records

1. Search or filter the Catalogue.
2. Select the intended records.
3. Review the selected IPs.
4. Choose the bulk delete action.
5. Confirm the count.
6. Verify active block rules separately.

72.7 Search by ASN

1. Open the ASN or organisation filter.
2. Enter the ASN number or organisation.
3. Review all matching IPs.
4. Compare event and PPC histories.
5. Do not bulk block the ASN solely because many records are visible.

72.8 Search by CIDR

Where CIDR search is supported:

1. Enter the range.
2. Verify the calculated first and last address where the UI shows it.
3. review the matched count.
4. export before bulk deletion.
5. process only the intended data category.

72.9 Enrich an IP

An enrichment action may retrieve:

- ASN;
- organisation;
- location;
- network type;
- abuse contact;
- reverse DNS.

Enrichment changes the information attached to the record. It does not itself block the source.

72.10 Add a manual note

A useful note includes:

- reason;
- supporting evidence;
- date;
- administrator;
- follow-up action;
- review date.

Avoid unsupported identity claims.

---

73. Blocking, unblocking, and range management

Blocking is an enforcement action. It is different from deleting a log.

73.1 Block one IPv4 address

1. Open **Access Control**, **Blocked IPs**, or the IP record.
2. Choose **Add block**.
3. Enter the complete IPv4 address.
4. Add a reason.
5. Choose a duration where available.
6. Confirm the administrator IP is different.
7. Save.
8. Verify the rule appears in the active list.
9. Verify the server rule where used.

73.2 Block one IPv6 address

Use the complete normalised IPv6 address or the interface's supported format.

Do not shorten or expand it manually if the system normalises addresses internally.

73.3 Block an IPv4 CIDR range

Example:

```
203.0.113.0/24
```

Procedure:

1. Open **Access Control -> CIDR or Range Rules**.
2. Select **Add range**.
3. Enter the CIDR.
4. Review the displayed network size.
5. Review matched known IPs.
6. Confirm that the range does not include the administrator.
7. Add a reason and evidence reference.
8. Prefer a temporary review period for broad ranges.
9. save.
10. Verify the active rule and server rule.

73.4 Block an explicit start and end range

Where start/end ranges are supported:

```
203.0.113.10
203.0.113.25
```

The product should convert the range safely or store it in a range-aware structure.

Do not approximate it with a larger CIDR unless the added addresses are intentionally included.

73.5 Remove one block

1. Search the active block list.
2. Select the address.
3. choose **Remove block**.
4. confirm database removal.
5. confirm server-rule removal.
6. retest.

73.6 Remove a CIDR block

1. Open range rules.
2. Search the exact CIDR.
3. Select the range.
4. choose **Remove**.
5. confirm the number of addresses affected.
6. verify `.htaccess`, Nginx, or external firewall.
7. clear applicable caches.
8. retest a known address in the range.

73.7 Reduce a range instead of removing it

Example: replace a `/24` with a narrower `/28`.

Procedure:

1. export or record the existing rule;
2. calculate the narrower range;
3. add the new narrower rule;
4. verify it;
5. remove the broad rule;
6. retest;
7. document the reason.

This order avoids an unintended gap during the change.

73.8 Temporary versus permanent blocks

Use temporary blocks for:

- uncertain patterns;
- short attacks;
- investigation;
- new ranges;
- campaign anomalies.

Use persistent blocks for:

- confirmed hostile sources;
- repeated exploit probing;
- manually reviewed persistent abuse.

73.9 Block source versus block route

Sometimes blocking a route is safer than blocking a network.

Examples:

- one junk path requested by many unrelated sources;
- parameter spam;
- obsolete attack route;
- unwanted endpoint.

Choose the narrowest effective control.

---

74. Whitelist operations

A whitelist is a powerful bypass and must be managed carefully.

74.1 Whitelist one IP

1. Open **Access Control -> Whitelist**.
2. Select **Add IP**.
3. enter the complete address.
4. add a reason.
5. add an expiry where appropriate.
6. save.
7. confirm the rule appears above conflicting block rules.

74.2 Whitelist the administrator

Before enforcement:

1. confirm the detected public IP;
2. compare it with an independent public-IP check;
3. confirm proxy configuration;
4. add it to the whitelist;
5. log out and back in;
6. test access.

74.3 Whitelist a CIDR range

Use only when every address in the range should receive the bypass.

Suitable examples may include:

- a tightly controlled office network;
- a trusted monitoring range;
- a documented provider callback range.

Avoid broadly whitelisting:

- an entire mobile carrier;
- all Cloudflare proxies as visitors;
- a large hosting provider;
- a public VPN range.

74.4 Remove a whitelist entry

1. Search the whitelist.
2. confirm why it was added.
3. review recent activity.
4. remove the entry.
5. monitor subsequent requests in Observation Mode.
6. do not automatically add a block.

74.5 Whitelist precedence

The product should document whether a whitelist overrides:

- IP blacklist;
- CIDR blacklist;
- ASN block;
- scoring;
- query verdict;
- 410 rule;
- login lock.

Absolute administrator safety and explicit whitelist behaviour must be predictable.

74.6 Time-limited whitelist

A temporary whitelist is useful for:

- developer access;
- migration;
- support;
- monitoring;
- testing.

The interface should show the expiry clearly.

---

75. Deleting an IP range from logs and records

Removing every trace of a range is a multi-area operation.

Deleting a range from the Activity Log alone does not necessarily remove it from the Catalogue, PPC Integrity, Evidence Ledger, block rules, or `.htaccess`.

75.1 Decide what the user actually wants

Goal A: stop blocking the range

Remove enforcement only. Preserve history.

Goal B: remove old clutter

Delete historical Activity Log data but preserve security decisions.

Goal C: satisfy a privacy or data-deletion request

Delete applicable personal data across all relevant tables while preserving lawful audit information where required.

Goal D: completely purge a test range

Remove test records, classifications, blocks, PPC records, evidence records, and generated rules.

75.2 Recommended full range cleanup order

Use this order for a complete purge:

1. Export relevant evidence.
2. Remove active block rules.
3. Remove whitelist entries if the range was trusted.
4. remove generated server rules.
5. remove Query Intelligence overrides tied to the range where applicable.
6. delete Activity Log records.
7. delete IP Catalogue records.
8. delete PPC records if intended.
9. delete Evidence Ledger records if intended and lawful.
10. clear cached dashboard aggregates.
11. run data reconciliation.
12. verify the range no longer appears.
13. retest access and logging.

Removing enforcement before deleting history ensures the administrator can still identify the active rule while the evidence exists.

75.3 Delete a CIDR from the Activity Log

1. Open **Activity Log**.
2. set the desired date range.
3. enter the exact CIDR in the range filter.
4. confirm the result count.
5. export.

6. select all matching records.
7. choose **Delete records**.
8. confirm.
9. refresh the view.

Fallback where CIDR filtering is absent:

- filter known addresses individually;
- use a verified exported list of IPs in the CIDR;
- process in batches;
- do not use an imprecise prefix for IPv6.

75.4 Delete the same CIDR from the IP Catalogue

1. Open **IP Catalogue**.
2. enter the exact CIDR.
3. review matched addresses.
4. select all matching Catalogue records.
5. choose **Delete records**.
6. confirm that active blocks were already removed.
7. run Catalogue reconciliation where available.

75.5 Delete the same CIDR from PPC Integrity

Only delete PPC records when there is a valid reason.

1. Open **PPC Integrity**.
2. select the provider or **All providers**.
3. set the date range.
4. filter by CIDR, IP list, ASN, or organisation.
5. export the evidence.
6. choose the data-management or delete action.
7. confirm whether the deletion affects:
    - arrival rows;
    - monthly summaries;
    - repeat-IP calculations;
    - evidence packs;
    - complaint history.
8. regenerate summaries.

Do not delete suspicious traffic merely because it makes campaign performance look poor.

75.6 Delete the same CIDR from the Evidence Ledger

1. Open **Evidence Ledger**.
2. filter by CIDR and date.
3. confirm event types.
4. export the original ledger.
5. select all matching rows.
6. choose **Delete**.
7. record the reason for the deletion.
8. regenerate reports.

Evidence deletion may have legal, contractual, or compliance consequences. Access should be restricted to authorised administrators.

75.7 Remove the CIDR from block rules

1. Open **Access Control**.
2. search the exact CIDR.
3. remove the block.
4. confirm server-rule deletion.
5. verify the range is not covered by a broader parent CIDR.

Example:

Removing:

```
203.0.113.0/28
```

does not restore access if this broader rule remains:

```
203.0.113.0/24
```

75.8 Remove child addresses covered by a parent block

If one IP was unblocked but the parent range remains blocked, either:

- remove the parent range; or
- add a specific whitelist exception if the rule engine and server support the required precedence.

Document the reason for the exception.

75.9 Delete an explicit IP list

Where the user has a list:

1. normalise and validate every address;
2. remove duplicates;
3. separate IPv4 and IPv6;
4. preview matched records;
5. export;
6. apply the action to the selected data area;
7. verify results.

The interface should support pasting one address per line.

75.10 Verify complete removal

Search the range in:

- Activity Log;
- IP Catalogue;
- Access Control;
- PPC Integrity;
- Evidence Ledger;
- Query Intelligence notes or overrides;
- complaint history;
- `.htaccess`;
- Nginx or external firewall;
- exports stored on the server;
- object cache.

A purge is not complete until all intended areas have been checked.

---

76. Query Intelligence operations

Query Intelligence manages decisions about URL parameters and query patterns.

76.1 Open Query Intelligence

1. Open **BotExorcist -> Query Intelligence**.
2. select the date range or analysis window.
3. review pending, blocked, approved, and overridden items.

76.2 Review a parameter

1. open the parameter or query record;
2. review example URLs;
3. review the route;
4. review frequency;
5. review source IPs and ASNs;
6. check whether administrators generated the parameter;
7. check campaign and analytics use;
8. choose a verdict.

76.3 Mark “Do not block”

Use when the parameter is legitimate.

Examples:

- `utm_source`;
- `gclid`;
- product filters;
- language parameters;
- legitimate form tokens.

The verdict should preserve history.

76.4 Mark “Block”

Use when the parameter is clearly unwanted or dangerous.

The action should state whether it:

- blocks the whole request;
- strips the parameter;
- returns 410;
- logs only;
- applies to one route or all routes.

76.5 Remove an override

1. open the parameter;
2. review the prior reason;
3. choose **Remove override** or return to automatic analysis;

4. preserve the audit history;
5. monitor future requests.

76.6 Delete query history

Deleting history should not silently remove the current verdict.

The UI should separate:

- delete observations;
- remove verdict;
- remove override;
- remove generated route rule.

76.7 Bulk approve tracking parameters

1. filter by a known tracking family;
2. review exact parameter names;
3. exclude suspicious lookalikes;
4. select approved parameters;
5. choose **Do not block**;
6. record the reason.

Do not approve every parameter beginning with `utm` if arbitrary unknown keys are being abused.

76.8 Reanalyse

A reanalysis action should:

- preserve manual overrides;
- recalculate automatic recommendations;
- record the analysis time;
- identify the model or rule version;
- avoid changing administrator-generated records.

---

77. PPC Integrity operations

77.1 Select the provider

Use the provider tabs:

- Google;
- Microsoft/Bing;
- Meta;
- All providers.

Provider selection should filter dashboard cards, tables, exports, and date comparisons consistently.

77.2 Change the date range

1. choose a predefined period or custom range;
2. confirm the site time zone;
3. apply;
4. wait for all cards and tables to refresh;
5. verify the displayed range.

A spinner that never finishes is a diagnostic problem, not evidence that no traffic exists.

77.3 Inspect a repeat IP

1. open **Repeat IPs**;
2. select the address;
3. review arrival dates;
4. compare campaigns;
5. review landing pages;
6. review human signals;
7. review ASN and organisation;
8. open the full IP Catalogue record.

77.4 Inspect a repeat ASN

1. open **Repeat ASNs**;
2. select the ASN;
3. review all associated IPs;
4. compare behaviour;
5. identify whether the network is:
    - mobile;
    - cloud;
    - hosting;
    - corporate;
    - residential.
6. avoid network-wide action without strong evidence.

77.5 Exclude the administrator from PPC analysis

Administrator and test traffic should be excluded or clearly labelled.

Use:

- administrator IP;
- authenticated administrator state;
- testing marker;
- manual exclusion.

Do not delete test traffic if preserving a complete audit is preferable. Classification may be better.

77.6 Delete test PPC records

1. filter by the test IP;
2. filter the exact testing period;
3. export if required;
4. choose **Delete test records** or the relevant bulk action;
5. regenerate cards and monthly summaries;
6. verify the Evidence Ledger separately.

77.7 Build an evidence pack

1. select provider;
2. set date range;
3. apply campaign or network filters;
4. select evidence categories;
5. add a factual summary;
6. include limitations;
7. generate the report;
8. review before sending.

77.8 Reset PPC dashboard cards

A card reset should not delete raw PPC evidence unless the confirmation explicitly says so.

Possible reset scopes:

- refresh calculations;
- clear cached cards;
- rebuild monthly aggregates;
- delete raw arrivals;
- delete all PPC evidence.

These actions must not share an ambiguous **Reset** label.

77.9 Correct an incorrect provider classification

1. open the arrival record;
2. review referrer and campaign parameters;
3. choose the correct provider or **Unclassified**;
4. record the reason;
5. rebuild provider totals.

---

78. Evidence Ledger operations

78.1 Search the ledger

Filter by:

- date;
- IP;
- CIDR;
- ASN;
- provider;
- event;
- action;
- campaign;
- route;
- block state.

78.2 Open one evidence event

Review:

- raw fields;
- derived fields;
- score;
- classification;
- administrator notes;
- source module;
- associated activity;
- associated IP record.

78.3 Add a note

A note should identify:

- author;
- date;
- purpose;
- evidence source;
- factual versus interpretive content.

78.4 Export CSV

Use CSV when:

- raw data is needed;
- another analysis system will process the data;
- field-level preservation matters.

78.5 Export spreadsheet

Use spreadsheet export when:

- sorting and filtering are needed;
- a client-facing workbook is appropriate;
- summary tabs are included.

78.6 Export PDF

Use PDF when:

- a fixed report is needed;
- evidence must be presented;
- the date range and filters should remain visible.

Keep the raw CSV as the source record.

78.7 Delete one ledger event

1. open the event;
2. confirm it is the correct record;
3. export or note it;
4. select **Delete**;
5. add a deletion reason where supported;
6. confirm.

78.8 Delete a filtered group

1. apply all filters;
2. review the count;
3. export;
4. select all matching;
5. delete;
6. rebuild summaries;
7. record the deletion reason.

78.9 Preserve complaint references

Do not delete evidence referenced by an active complaint without recording:

- complaint identifier;
- recipient;
- submission date;
- reason for deletion;
- replacement evidence.

---

79. Login Shield operations

79.1 Enable successful-login trust

1. open **Login Shield**;
2. enable trust after a correct login;
3. choose whether trust applies to:
   - session;
   - IP;
   - ASN;
   - a limited duration.
4. save;
5. test login from the normal administrator network.

79.2 Disable strict IP lock

Use when the administrator frequently changes networks.

1. open Login Shield settings;
2. disable strict IP locking;
3. keep authenticated-session trust enabled where appropriate;
4. test mobile and desktop login.

79.3 Enable ASN locking

ASN locking can tolerate changing IPs within one network but may create risk on large or shared networks.

Before enabling:

- identify the administrator ASN;
- confirm it is not a large shared mobile carrier;
- understand travel and VPN implications;
- preserve emergency recovery access.

79.4 Remove a trusted login record

1. open trusted sessions or login trust;
2. locate the user, IP, or ASN;
3. revoke trust;
4. do not automatically block;
5. require a fresh successful login.

79.5 Review a failed-login source

1. open failed login activity;
2. review username attempts;
3. review frequency;
4. review IP and ASN;
5. review whether the source also probed routes;
6. choose observe, temporary block, persistent block, or no action.

79.6 Recover from login lock

Use section 41 or 64.

After recovery:

- disable strict locking;
- confirm IP detection;
- remove the accidental block;
- whitelist the correct administrator;
- test logout and login.

---

80. Server rules and .htaccess operations

80.1 View managed rules

Where the plugin exposes a rule viewer:

1. open **Access Control** or **Server Rules**;
2. select the managed-rule view;
3. review allow rules;
4. review deny rules;
5. confirm whitelists appear before blacklists.

80.2 Back up `.htaccess`

Before editing:

1. open the WordPress root;
2. download `.htaccess`;
3. name the backup with a date;
4. store it outside the web root where possible.

80.3 Remove one generated IP rule manually

Use only when dashboard access is unavailable or the rule is stuck.

1. find the BotExorcist-managed section;
2. identify the exact IP;
3. remove only that rule;
4. preserve section markers;
5. save;
6. test the server;
7. remove the corresponding database block after login.

80.4 Remove one generated CIDR rule manually

1. identify the exact CIDR or generated expression;
2. confirm it belongs to BotExorcist;
3. remove the line or complete generated block entry;
4. save;
5. test;
6. remove the matching database rule.

80.5 Do not delete WordPress rewrite rules

The standard WordPress section often includes:

```
# BEGIN WordPress
...
# END WordPress
```

Do not delete it while removing a BotExorcist rule.

80.6 Rebuild managed rules

A rebuild action should:

- back up the current file;
- preserve unrelated rules;
- write whitelist entries first;
- write blacklist entries after;
- validate syntax;
- log success or failure.

80.7 Reconcile database and server rules

Use reconciliation when:

- the dashboard says an IP is blocked but `.htaccess` does not;
- `.htaccess` contains a rule absent from the dashboard;
- a restore reverted one side;
- manual editing occurred.

The reconciliation screen should show:

- database-only rules;
- server-only rules;
- matching rules;
- invalid rules;
- recommended actions.

80.8 Nginx and other servers

Do not present `.htaccess` instructions as universal.

For Nginx or managed firewalls, document the actual rule location or use application-level blocking when server-rule integration is unavailable.

---

81. Bulk deletion, retention, and privacy requests

81.1 Retention before deletion

A retention policy is safer than repeated manual clearing.

Configure separate retention where appropriate for:

- Activity Log;
- IP Catalogue;
- PPC records;
- Evidence Ledger;
- complaint history;
- diagnostics;
- cached aggregates.

81.2 Delete records older than a date

1. open the relevant data-management screen;
2. select **Older than**;
3. choose the cutoff;
4. preview the count;
5. export if required;
6. delete;
7. optimise or clean tables where supported.

81.3 Keep security evidence longer than display data

A practical policy may keep:

- detailed request logs for a shorter period;
- security incidents longer;
- complaint evidence longer;
- aggregate statistics longest.

The actual policy must match privacy obligations and business needs.

81.4 Privacy deletion request by IP

An IP address can change and can be shared, so an IP-only request may not prove identity.

Where the website owner lawfully decides to process the request:

1. verify the scope;
2. identify the date range;
3. search all data areas;
4. export an internal deletion record;
5. delete applicable records;
6. preserve legally required information;
7. record completion.

81.5 Bulk deletion preview

Before confirmation, show:

- selected filter;
- tables affected;
- records affected;
- active rules affected;
- evidence affected;
- whether summaries will rebuild;
- whether the action is reversible.

81.6 Soft delete versus permanent delete

A soft delete:

- hides the record;
- allows recovery;
- preserves an audit entry.

A permanent delete:

- removes the record;
- may be irreversible;
- may affect reports.

The interface must identify which action is being used.

81.7 Scheduled cleanup

Scheduled cleanup should:

- run in batches;
- avoid long locks;
- preserve manual classifications where configured;
- log completion;
- report failures;
- not delete active rules by accident.

---

82. Reset controls and what they remove

The word **Reset** is dangerous unless the scope is explicit.

82.1 Reset visible filters

Removes:

- search term;
- date filters;
- event filters;
- provider filters.

Does not delete data.

82.2 Refresh dashboard calculations

Recalculates cards from existing data.

Does not delete raw records.

82.3 Clear dashboard cache

Removes cached totals.

Does not delete the source records.

82.4 Reset Activity Log

May delete Activity Log rows.

Should not automatically delete:

- IP Catalogue;
- PPC records;
- Evidence Ledger;
- block rules;
- whitelists.

82.5 Reset PPC Integrity

The interface must distinguish:

- clear cached cards;
- rebuild summaries;
- delete PPC arrivals;
- delete all PPC evidence.

82.6 Reset IP Catalogue

May delete IP intelligence records.

Should state whether it also removes:

- manual notes;
- classifications;
- blocks;
- whitelists;
- enrichment;
- PPC links.

82.7 Factory reset

A factory reset may remove:

- settings;
- logs;
- Catalogue;
- PPC records;
- Evidence Ledger;
- rules;
- licence association;
- local caches.

It should require:

- administrator capability;
- nonce;
- typed confirmation;
- backup warning;
- exact list of affected data.

82.8 Never combine unblock with delete silently

An action should not remove both enforcement and forensic history unless the user explicitly chooses both.

---

83. Common administrator workflows

83.1 “This customer is blocked. Let them back in.”

1. Search the IP in Activity Log.
2. open the IP Catalogue record.
3. confirm the block reason.
4. remove the active block.
5. check parent CIDRs.
6. check `.htaccess`.
7. add a temporary whitelist if justified.
8. retest.
9. preserve the history.

83.2 “Delete our office testing traffic.”

1. identify office IP or CIDR;
2. confirm the test date range;
3. export if needed;
4. delete Activity Log rows;
5. delete PPC test rows;
6. delete Evidence Ledger test rows if appropriate;
7. preserve the office whitelist;
8. rebuild summaries.

83.3 “Remove a hostile range completely.”

1. export evidence;
2. decide whether history must be retained;
3. remove the active block only if access should be restored;
4. otherwise keep enforcement and delete only expired clutter;
5. use the full purge process only when genuinely required.

A hostile range usually should not be purged merely because it is hostile. Its history may be valuable.

83.4 “A legitimate query parameter is returning 410.”

1. open Query Intelligence;
2. search the parameter;
3. review affected routes;
4. select **Do not block**;
5. remove the 410 override;
6. clear CDN and page cache;
7. retest;
8. preserve the old decision history.

83.5 “Elementor or Divi cannot save.”

1. switch to Observation Mode;
2. confirm administrator trust;
3. reproduce the save;

4. find the exact blocked request;
5. identify AJAX or REST route;
6. add a narrow safe lane;
7. retest;
8. return to enforcement gradually.

83.6 “The dashboard is wrong after deleting records.”

1. refresh date range;
2. clear dashboard cache;
3. rebuild aggregates;
4. verify retention;
5. check that another table still contains the records;
6. confirm provider filter.

83.7 “Credits are zero.”

1. confirm the notice states the firewall remains active;
2. identify which optional service paused;
3. add credits or disable the optional service;
4. verify local firewall events continue;
5. contact support if the core firewall stopped.

83.8 “The same ASN appears repeatedly.”

1. open the ASN record;
2. review network type;
3. review all associated IPs;
4. compare routes and behaviour;
5. identify paid arrivals;
6. avoid whole-ASN blocking without strong evidence;
7. consider a complaint or narrow CIDR action.

---

84. Understanding action scope

Every destructive or enforcement action should display its scope.

84.1 Scope labels

Recommended labels:

- **View only**
- **History only**
- **Current table only**
- **All matching records**
- **Active enforcement**
- **Server rule**
- **PPC evidence**
- **All BotExorcist data**

84.2 Confirmation example: delete Activity Log range

```
You are about to permanently delete 4,218 Activity Log records
matching 203.0.113.0/24 from 1 July to 15 July 2026.

This will not remove:
- IP Catalogue records
- PPC records
- Evidence Ledger records
- active block rules
- .htaccess rules

Type DELETE to continue.
```

84.3 Confirmation example: remove CIDR block

```
You are about to remove the active block for 203.0.113.0/24.

This will:
- remove the database block rule
- remove the generated server rule where present
- restore access for addresses in the range unless another rule applies

Historical logs will remain.
```

84.4 Confirmation example: complete purge

```
You are about to purge records for 203.0.113.0/24 from:
- Activity Log
- IP Catalogue
- PPC Integrity
- Evidence Ledger
- cached summaries
```

```
Active rules are handled separately.

This action cannot be undone after the backup window expires.
```

84.5 Result summary

After an operation, show:

- records matched;
- records changed;
- rules changed;
- failures;
- tables skipped;
- next recommended step.

---

85. Undo, rollback, and recovery

85.1 Undo window

Where practical, destructive actions should provide a short undo window using soft deletion.

85.2 Restore from export

A normal CSV export is not automatically a restorable backup.

A true restore package needs:

- stable identifiers;
- field mapping;
- schema version;
- relationship data;
- import validation.

85.3 Restore an accidental block removal

1. find the original rule in the audit log;
2. confirm the range;
3. re-add it;
4. restore the reason;
5. verify server rule;
6. monitor.

85.4 Restore an accidental whitelist removal

1. identify the trusted address;
2. confirm it is still safe;
3. re-add the whitelist;
4. add an expiry if appropriate;
5. test.

85.5 Recover deleted logs

Recovery may require:

- soft-delete restore;
- database backup;
- hosting snapshot;
- external archive.

If no backup exists and permanent deletion completed, recovery may be impossible.

85.6 Audit administrative changes

Record:

- administrator;
- time;

- action;
- target;
- previous value;
- new value;
- reason;
- result.

85.7 Emergency rollback after an update

1. preserve diagnostics;
2. disable enforcement;
3. restore previous plugin ZIP;
4. check schema compatibility;
5. restore server rules if required;
6. test access;
7. report the regression.

---

86. Operational terminology reference

Active block
A currently enforced deny rule.

Aggregate
A calculated summary derived from raw records.

Bulk action
An operation applied to multiple selected or matching records.

Catalogue record
A persistent intelligence record for an IP address.

Clear filters
Returns a view to its default filters without deleting data.

Complete purge
Removal of all selected data categories for a target.

CIDR filter
A mathematically range-aware search for addresses within a network.

Delete history
Removes stored records but does not necessarily change access.

Enforcement rule
A rule that actively allows, blocks, or otherwise controls requests.

Parent CIDR
A broader network containing a narrower IP or child CIDR.

Rebuild aggregates
Recalculates dashboard summaries from stored source records.

Reconcile
Compares two sources of truth, such as database and `.htaccess`, and resolves differences.

Remove block
Stops active enforcement while ordinarily preserving history.

Reset
An ambiguous term that must always be qualified by scope.

Soft delete
Marks a record deleted while allowing recovery.

Whitelist exception
A narrow allow rule that overrides a broader deny rule where precedence supports it.

---

87. Final product statement

BotExorcist gives the website owner an independent record of:

- who arrived;
- what was requested;
- how the visitor behaved;
- which network was involved;
- which human signals were observed;
- which suspicious events occurred;
- how risk changed;
- why the system responded.

BotExorcist does not replace human judgement.

It gives human judgement better evidence.
