BROKER CONSOLE · MODULE REFERENCE

Trade Monitoring

Who is trading against the book rather than in it: the risk desk, the abuse detectors, the fraud cases a person decides on, the limits that hold them, and the credit the brokerage put on an account.

28 detectors, five families

Evidence, or no action

No pre-trade hook — said on every screen

Reactions built from what the gateway admits

MT4 declared absent, never zeroed

PRODUCT	MODULE	CONTENTS	VERSION
HonorPro Hub	Trade Monitoring	5 menus · 22 tabs	1.0 · 24 Sep 2026

Risk & Abuse	Dashboard · Device ID · IP Address · Arbitrage Detection · Scalpers Detection · Swap Abuser · Dividend Abuser
Anti Fraud	Overview · Cases · Money Flow · Toxic Flow · Detectors & Settings · Scans · Statistics · Journal
Anti-Abuse Protection	Protection Rules — the flags, their evidence, and the measures taken
Limitation	Exposure Limit (MT5) · Exposure Limit (MT4) · Anti Scalp · Auto-Hedger
Bonus & Credit Manager	Bonus & Credit · TradeShield

01 THE MODULE

Detection after the fact, stated on every screen

One property of this platform decides the shape of every screen in this module, and a broker should know it before buying: the plugin this installation runs registers five sinks with the trading server — tick, user, position, order and deal — and every one of them fires **after** a deal is booked. There is no request sink, so nothing here refuses an order at the moment it is placed.

5
MENUS
Risk & Abuse, Anti Fraud, Anti-Abuse, Limitation, Bonus & Credit

22
TABS
Each one an address a person can paste to a colleague

28
DETECTORS
Five declare themselves unavailable, each with its reason

What the platform can actually do

Exposure Limit is monitoring plus reaction; Anti Scalp is detection plus reaction. Both run after the breaching trade has been booked — usually within seconds, never before. The whole set of reactions is:

- **notify** — the desk is told through the notification component, once per breach rather than once per sweep
- **close to the ceiling** — largest position first, stopping at the ceiling rather than flattening the subject
- **close everything** — only where the limit's subject *is* the account, the one scope on which "everything" is what the limit meant
- **withdraw the trading permission** — the account keeps what it holds and can open nothing further

One book, read once

The exposure aggregation is computed on the gateway in one pass and cached; the Risk Dashboard, Exposure Limit and the Auto-Hedger all read that one snapshot rather than each making a pass of their own.

The gateway **refuses to store** a reaction of "refuse" or "delay" and says why: refusing a close needs a pre-trade request hook inside the trading server, and delaying one needs the same hook held open. Every reaction picker in this module is built from what the gateway admits it can run, never from a list typed into a screen.

MT4 is not on this gateway

This gateway speaks MT5 only. It has no MT4 connection, no MT4 book and no MT4 symbol table. The MT4 tabs say so in the gateway's own words and take **no risk snapshot at all** — reading the MT5 book and labelling the answer MT4 is the one mistake this module must not make. A limit set written for MT4 is stored and marked unevaluated.

No IP, no device per login

Neither backend exposes one. The shared-identity detector declares itself unavailable in as many words, and the screens that ask the question print that sentence instead of an empty table.

02 MENU · RISK & ABUSE

The book right now, and the two questions this gateway cannot answer

Seven tabs: one live risk board, four narrow questions put to the fraud engine's own families, and two abuses a brokerage loses real money to, each given a desk. The first three are here; the rest overleaf.

Dashboard `/risk`

The whole book, added up on the gateway in a single pass and cached for a few seconds, so four panels refreshing together cost the broker's machine one read. Exposure by instrument and by asset class, the margin read, the client leaderboards, every breach of the desk's own limit rules, and a statistics catalogue over a chosen period beside the one before it.

Three things it refuses to smooth over. **Sign is direction** — net is long minus short, so a short book is drawn negative. **Margin is not the whole book** — MT5 has no book-wide margin read, so the gateway reads the most exposed accounts up to a cap and the screen says "the N most exposed", never "the book". And **a symbol with no contract size is counted in lots and not valued**, with the reason above the table.

Overview · Clients · Breaches · Statistics · Limit rules · One pass, cached

Device ID `/fraud/section/coordination?view=device`

Which accounts are being traded from one machine. The screen reads the Coordination family's own evidence looking for a device on each case and groups by whatever it finds. On this gateway it finds nothing, and it says so in the gateway's words rather than drawing an empty table: the shared IP and device detector is published permanently unavailable because **neither the MT5 manager backend nor the plugin bridge exposes login records** — only the last access time on the account.

Coordination family · Grouped by device · Declared unavailable, with the reason

IP Address `/fraud/section/coordination?view=ip`

The same question asked of an address: several logins from one address is the first thing a desk asks for when it suspects one person behind a book of clients. The same answer, and the same honesty about it. What is shown instead is the identity evidence this gateway does serve — the same e-mail, phone, name and country, or the same private mail domain behind several logins. The day a release puts an address on the evidence, the same code groups by it without being rewritten.

Coordination family · Grouped by address · Identity signals shown instead

The route that **does** work is in Server Management: the Agent on the broker's own machine reads the login lines out of the server's log directory, where the server itself writes them. That is the IP Geolocator, and it works identically on MT4. Both tabs here keep their place because the question is a real one and the day a release puts an address or a device on a case's evidence, these screens group by it unchanged.

03 MENU · RISK & ABUSE

Four desks for the four arguments a dealer has

Two fraud families given their own tab because they are the two a dealing desk opens daily, and two abuses given a desk of their own because the evidence they need does not fit a generic case row.

Arbitrage Detection `/fraud/section/arbitrage`

Arbitrage and hedging across accounts: opposite positions held on several logins, bonus hedging, swap-free carry, dividend abuse and correlated-pair latency. Five detectors, each with a paragraph of what it looks for, the values in force, and the queue the family raised.

5 detectors · Cases with evidence · Values in force shown

Scalpers Detection `/fraud/section/scalping`

Scalping and latency: holding-time scalping, latency arbitrage, rollover-hour scalping, pip hunting, stop hunting and spread-widening. This is the statistical half of the subject — median holding time, trades per active hour, win rate, a score and a case. Two of the six declare themselves unavailable and say why: the deal feed carries no stop-loss or take-profit levels, and this gateway keeps no tick or spread history.

6 detectors · 2 unavailable, with reasons · Anti Scalp is the broker's own rule

Swap Abuser `/swap/abuse`

The gateway's anti-abuse desk as a ranked table, worst first: flags raised over closed trades, each with a score, a severity, a status and the measures taken on it. The actions are the gateway's own vocabulary — review or dismiss a flag, apply a measure, lift one — and a scan is a real job, polled until it is done. Under the figures sits that desk's own statistics catalogue, and a breakdown row drills straight into the flags behind it.

Ranked flags · Scan as a job · Statistics with drill-down

Dividend Abuser `/fraud/dividend-abuser`

One detector given a desk of its own — it is declared in the same registry, in the Arbitrage family, tuned through the same settings document and raising cases into the same queue. What it gets is the shape of the question: under a chosen case, the **actual positions with their ex-dates and the dividends credited**, which is three tables' worth of fact a generic case row cannot show; and beside the thresholds, a live what-if — with these values, N accounts would be flagged now against M under the values in force.

It reads the Dividend Manager's ledger against the position history. When the calendar is empty the gateway says so in the detector's own reason: it is *available and finds nothing*, which is not the same as unavailable, and the screen offers the door to the Dividend Manager rather than drawing a nought.

Cases · Thresholds · Live what-if · Dividend Manager ledger

A family tab lists its detectors with the values in force and then the queue that family raised. A detector this installation cannot run **keeps its description and its parameters' place** and says why it cannot run, rather than disappearing from the list.

04 MENU · ANTI FRAUD

Every kind of illegitimate trading, as cases a person decides on

Twenty-eight detectors across five families — scalping, arbitrage, coordination, money flow and toxic flow — raising cases into one queue with a severity, a score, evidence and a status. The queue is the product; the detectors are how it fills.

Overview `/fraud`

The desk at a glance, composed from five gateway doors each in its own envelope: the policy in force, the last scan, the open cases by severity, the score distribution and the money at risk, and one card per family with its open count. A door that did not answer names itself on its own card while the others draw — a desk is not blanked because one read failed.

Policy in force · Open by severity · Money at risk · Per-family counts

Cases `/fraud/cases`

The queue, worst first, with every filter the gateway's list takes carried in the address — status, severity, family, detector, login, assignee, tag and free text — so a queue somebody is arguing about is a queue they can paste to a colleague. A row opens the case drawer; ticking rows arms the bulk bar, and every bulk action is confirmed with a sentence naming how many cases and what happens to them. Dismissing or closing needs a reason, in bulk exactly as one at a time.

Filters in the address · Case drawer · Bulk with a named count · Dismissal needs a reason

Money Flow `/fraud/section/money`

Deposit-to-withdraw churn, bonus abuse, rapid cycling, chargeback signals and the first-deposit maximum-leverage burst. Five detectors; the chargeback one declares itself unavailable because the payments ledger carries no chargeback or reversal state.

5 detectors · 1 unavailable

Toxic Flow `/fraud/section/toxic`

Flow that is legitimate and still costs the book: short-horizon winners against the house, order flooding, martingale and grid blow-ups, news-only and off-hours trading, and volume far outside an account's own history. Order flooding declares itself unavailable — no order placement, modification or cancellation log exists for it to read.

6 detectors · 1 unavailable

The **gateway owns the vocabulary of families**. The five above keep the labels and the order this console gives them; a gateway that reports a sixth this console has never heard of grows a tab of its own beside them, in the gateway's order. A gateway that reports only the five changes nothing on screen.

05 MENU · ANTI FRAUD

The rules, the runs, the figures and the trail

The four tabs that are not a queue: what the engine is configured to look for, the jobs that make it look, what it found over a period, and every step somebody took.

Detectors & Settings `/fraud/settings`

Every rule the engine runs, editable. Each detector carries its parameters, a weight, severity weights, a whitelist and a live preview of what it would catch. The rail is the settings document's own contents — policy, scoring, the detectors under their family heading, the global whitelist, the scan schedule and the kept versions — and the save bar sits above it, because a draft spans the whole document and nothing unsaved may hide behind a closed pane.

Two things here are deliberately slow. Arming **auto** mode — where a measure fires at a score without a person in the loop — asks for a confirmation stating what auto would do on the next sweep with the draft values. And a parameter outside the gateway's range is refused by the gateway, with its sentence shown verbatim, rather than clamped quietly.

Policy: advise or auto · Scoring · Per-detector thresholds · Whitelists · Versions with reasons

Scans `/fraud/scans`

A scan is a real gateway job: a window of days, the detectors it ran, a cap with a flag saying the read was truncated, and a report of how many accounts, deals and windows it actually read — so "no findings" can be told apart from "nothing was read". A dry run finds and reports without opening a case, and the badge says so on the row and on the button before it is pressed. A running scan is polled and can be cancelled.

Dry run · What was read, counted · Cancellable · Run history

Statistics `/fraud/stats`

Detection by detector, family, severity, group, country, instrument, hour and weekday; the review backlog and its precision; the measures taken; the flagged accounts' share of the book; money at risk; scan health. The gateway publishes a catalogue declaring each figure's kind, which direction is good and the read it came from, and the screen draws whatever the catalogue describes. A drill on a breakdown row opens the case list filtered to it.

Self-describing catalogue · Period compare · Drill to the cases

Journal `/fraud/journal`

Every step, by whom, newest first: settings changes with what changed, scans, case decisions with their notes, and measures applied and lifted. Filterable by kind, case and login, exportable as the rows on screen or through the gateway's own export.

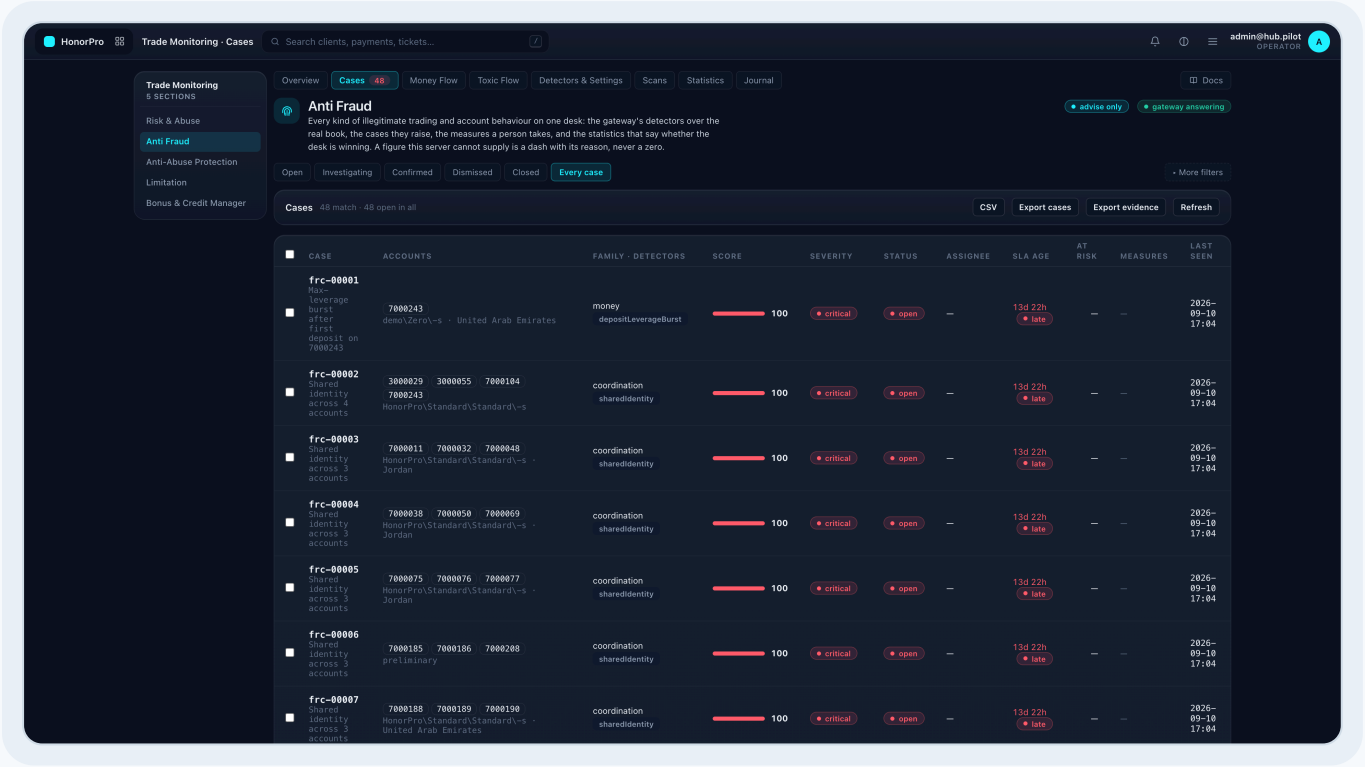
Settings · scans · decisions · measures · CSV export

Every save of the settings document is a numbered version carrying the operator's own reason, restorable from the same screen. **A reason is required**: the thresholds a detector ran under last quarter are part of the answer to why an account was flagged then.

06 THE FRAUD DESK

A queue where every row names what raised it

The Cases tab on a live installation. Every filter above the table is carried in the address, so a queue somebody is arguing about is a queue they can paste; the two pills at the top right say whether the gateway is answering and which policy is in force.



*Anti Fraud · Cases. A demo installation seeded with trading history, so these are the detectors' own findings rather than a brokerage's book. **advise only** means no measure fires without a person deciding.*

What a row carries

The case and what it is, the accounts it names with their group and country, the family and the detector that raised it, the score, the severity, the status, the assignee, how long it has been open against the desk's own service level, the money at risk and any measure already taken. Worst first, and a row opens the drawer with the evidence behind it.

Two exports, kept separate

One export is the table. Beside it is an export of the evidence, because the tickets and the thresholds behind an accusation are what has to travel with it — and it is large, which is why it is a second button rather than a column nobody asked for.

07 MENU · ANTI-ABUSE PROTECTION

An accusation you can show the client it is about

Every other screen in this console reports a state. This one makes an accusation, and an accusation that cannot be shown to the client is worth nothing — so the evidence, the metrics computed from it and the thresholds those metrics were compared against are all stored with the flag and served with it.

Protection Rules `/abuse`

The standing rules that watch the book without anybody asking, and what they caught. Five detectors run in batch over the closed-trade history: latency (short, almost always profitable trades clustered around news), scalping (a median holding time of seconds at a frequency nobody sustains), coordination (several accounts entering the same way within seconds), a news-window volume spike, and a hedge farmed across logins. Each has its own thresholds on the screen, and each is switched on or off there.

A flag opens with its evidence in tabs — the tickets behind it, the events, the flag's own history and the measures on it. "Score 87.6" is not a reason; "median holding time 4 seconds against a threshold of 30" is, and that is what the screen draws. The measures are the gateway's: monitor, flag, move the accounts into a restricted group, or cap their leverage. Both of the last two are reversible and are lifted from the same screen; a restrict is shown as a plan — these logins, into that group — before anything is pressed, and takes the accounts permission on top of this one, because it edits a client's account.

5 detectors · Evidence in tabs · monitor · flag · restrict · leverage cut · Scans · Journal

A status nobody types

A flag is open, reviewed, measured or dismissed, and **measured is reached only by applying a measure** — so the status and the account can never disagree. The buttons offered are the ones the gateway's own transition table allows from where the flag actually is; this console holds no copy of that table. Dismissing is the one move that refuses to happen without a sentence.

One scan at a time

A scan walks the whole closed-trade history in stages and reports its stage and percentage as it goes. One runs at a time for the whole gateway, and a second started while one is running is refused by name — the refusal is the sentence to show, rather than a queue nobody is watching.

Every measure lifts. A restrict records the group it moved the account out of and lifting it moves the account back there; a leverage cut is held by the leverage desk and lifted from here. An account whose leverage a person has changed by hand since is skipped rather than overwritten — the later human decision outranks the lift.

A repeat finding on an open or reviewed flag refreshes it; on a flag somebody has already decided it increments a recurrence count and leaves the evidence that decision was signed against exactly as it was. A seat scoped to a basket of groups sees only the flags naming one of its own logins and cannot start a scan at all.

The desk's own settings sit beside the detectors: how many days a scan reads back, the cap on trades read per login, how many tickets are kept as evidence on each flag, the score a firing detector starts at and the scores at which a flag becomes medium and then high, and the group a restrict moves an account into.

08 MENU · LIMITATION

What the book will not be allowed to carry

Three ceilings and one cover desk. The first three are after the fact and say so above everything else on the page; the fourth is the only thing in this module that sends an order of its own, and it never sends one unasked.

Exposure Limit (MT5) `/limits/exposure/mt5`

A ceiling per instrument, per group, per group-and-instrument and per account, measured on gross or net notional or gross or net lots — net on its absolute value, because a book 40 million short is as exposed as one 40 million long. Each ceiling carries a warn level, a reaction and an optional note, and the figures come from the same cached snapshot the Risk Dashboard reads, so opening this screen costs the broker no extra pass. A ceiling of zero is refused with the reason: it is in breach the moment anybody opens anything, and "not to be traded" is a group permission, not a limit.

Reacting is a two-step. The dry run draws the plan — which tickets, whose, how much it would shed and what would be left against the ceiling — and running it for real needs the word typed back. The plan stops at the ceiling rather than flattening the subject; where it cannot get there with the positions it can see, it says so.

Live · The Limit Set · Group × Symbol · Journal · Exempt logins, shown not skipped

Exposure Limit (MT4) `/limits/exposure/mt4`

The same document, the same four scopes and the same editor, against an MT4 server this installation does not have. No snapshot is taken at all: every figure is a dash with the gateway's sentence beside it, and the reaction picker is empty rather than offering choices that are all inert. The limits are stored, and start being evaluated the day an MT4 server is attached — without being rewritten.

Stored, not evaluated · No snapshot taken · No reactions offered

Anti Scalp `/limits/anti-scalp`

The broker's own floor on holding time, per group and per instrument — thirty seconds on the majors, five minutes on gold — with the exemptions the desk agreed in writing. Rules are tried in document order and the first enabled match wins, which cannot be read off a list of masks, so the editor carries a "which rule would govern this pair" check and the rows can be reordered. An account is listed once its rule's minimum trade count is met: one short trade is a client who changed their mind.

It is not a second detector. The Anti Fraud desk finds scalpers statistically; this is the term the brokerage sold. The screen reads that desk's open scalping cases and offers the door to them rather than competing with it.

Caught · The Policy · Journal · monitor · notify · restrict · Exemptions shown out loud

Firing a reaction takes the accounts permission **on top of** the desk's own, and the request has to carry the confirmation word: reading a limit set and closing a client's positions are not the same grant, and a refusal names both halves.

A holding-time floor above 24 hours is refused. That is not a scalping rule but a minimum holding period, and writing one here would catch every intraday trade on the book.

09 MENU · LIMITATION · AUTO-HEDGER

The one desk here that places an order

A brokerage holding the other side of its clients' book is short what they are long. The Auto-Hedger covers that net exposure on a coverage account, and it is one tab with seven screens inside it because it arrived as a product rather than as a page.

Auto-Hedger `/hedger`

A policy of rules in priority order, each matching instruments by name or by class, with a ladder of thresholds and the share of the exposure to cover at each rung, a hysteresis floor to unwind back through, trading sessions, lot rounding and caps. One rule governs an instrument — the first match wins and a shadowed rule says so. Only the shortfall is ordered, never the raw target, and the coverage accounts' own fills are excluded from the exposure the hedger reads, so it can never hedge itself.

Book · Rules · Statistics · Simulator · Incidents · Journal · Settings

Advise, and then live

The default mode is **advise**: every decision is computed, shown and journalled, and nothing is sent. Going live is a deliberate arming with its own confirmation. Above both sits a kill switch that is checked last, after every other reason a decision was skipped, so a desk can see what is being held back rather than an empty screen. Stopping it takes a reason; that reason is the only record of what the next person on the desk is walking into.

What it refuses

A hedge is checked against the coverage account's own free margin and a margin floor before it is sent, and refused in a sentence naming the numbers and ending "Nothing was sent". A refusal is an incident, because the exposure it did not cover is still open and somebody has to cover it by hand. An account it cannot price is not blocked — a hedger that refused every hedge it could not value would stop covering the book.

The seven screens

Book is the per-instrument table the desk lives in: client exposure, standing cover, target, residual, the band it is in and what the next sweep would do. **Rules** is the policy as a table, saved whole. **Simulator** runs a draft against the current book beside the saved policy's own answer, and over past sweeps where they were recorded. **Incidents** holds what was placed, refused or rejected, each acknowledged with a reason. **Journal**, **Statistics** and **Settings** complete it.

Looking is free; sending is not

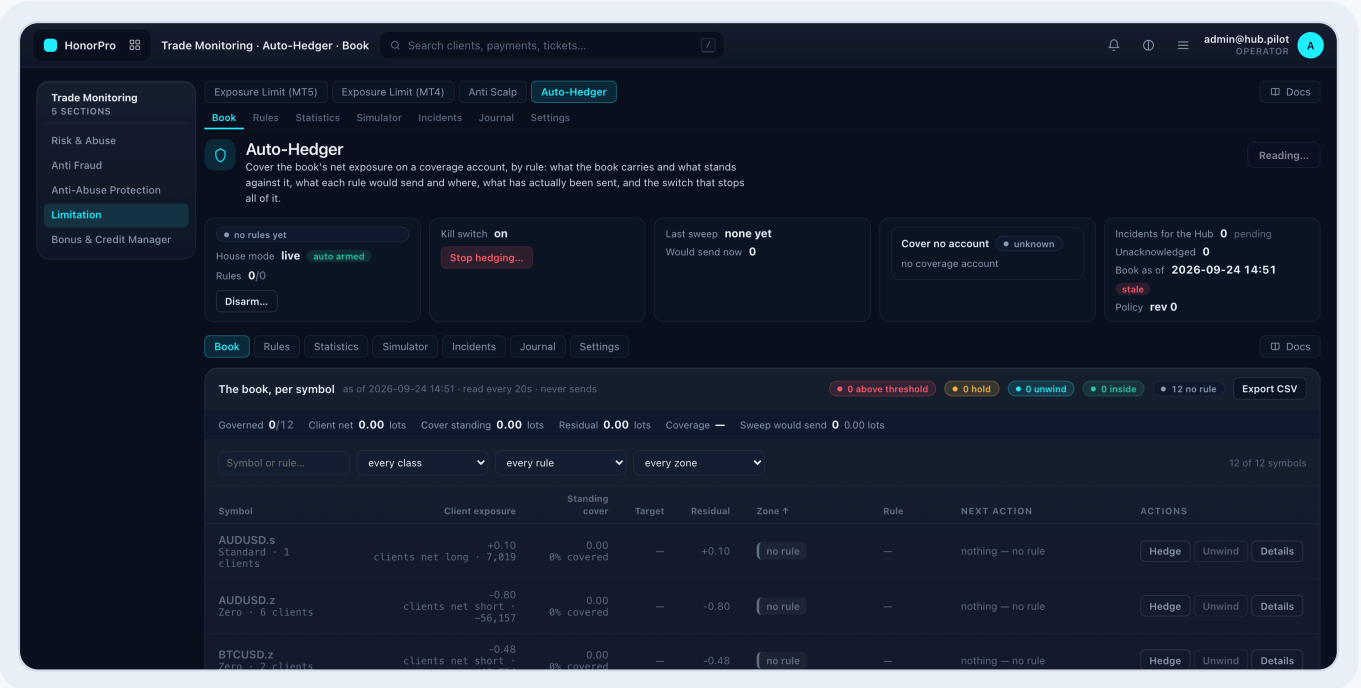
The screen polls status, the book and the preview. Running a sweep, hedging, unwinding, retrying, arming and the kill switch are the only calls that move money, every one sits behind a confirmation, and **no timer on this screen ever makes one**. A manual hedge and an automatic one are indistinguishable in the journal except for the rule name.

The sweep keeps no clock of its own inside the gateway: whatever drives it — this screen, a scheduler, the Agent — calls the same door under the same guarantees. Under a seat scoped to a basket of groups the desk all but closes, and only the preview survives, computed over that operator's own basket with the coverage account blanked.

10 THE HEDGING DESK

What the book carries, and what stands against it

The Auto-Hedger's Book tab. The strip across the top is the desk's whole state in five readings — the policy, the kill switch, the last sweep, the cover account and the incidents — and under it the per-instrument table a dealer lives in.



*Limitation · Auto-Hedger · Book. This installation has been armed live but carries no rule and names no coverage account, so every one of the twelve instruments on the book reads **no rule** and the sweep would send nothing.*

What a row carries

Per instrument: the clients' net exposure with the side spelt out, the cover standing against it and the share of the exposure it covers, the target, the residual, the band the row is in, the rule that governs it and what the next sweep would do. Hedge and Unwind open a preview of the order on the cover account, and nothing is sent until it is confirmed.

The state strip

"Read every 20s · never sends" is the screen saying what it is: the poller asks for the status, the book and the preview, and no timer here places an order. A book read that has gone stale is marked **stale** beside its timestamp rather than presented as current.

11 MENU · BONUS & CREDIT MANAGER

Credit put on an account, and what takes it back

Both tabs book a credit deal on a client's account, so both need the money permission on top of their own, and both are preview-first: there is no button on either screen that pays without having drawn the plan.

Bonus & Credit `/bonuses`

Campaigns — percentage, fixed or welcome — each a rulebook: the deposit it applies to, the caps per grant and per client, the budget, the dates, and how much volume turns the credit into the client's own money. A grant freezes the campaign's terms as they were the day it was made, so a later edit never rewrites what a client was promised.

Four jobs on four tabs, because different people do them at different moments. **Grants** is the ledger and the one button that converts credit to balance. **Sweeps** converts the eligible and expires the overdue across a campaign — a real run is not offered until a dry run has drawn the plan. **Withdrawals** previews what a payout costs the bonus; applying it books one credit deal, recording the money desk's decision rather than the payout. **Report** and **Journal** say where the money went and what was refused.

Grants · Sweeps · Withdrawals · Report · Journal · Idempotency key per form

TradeShield `/tradeshield`

The brokerage refunds a new client's early losses as credit. A rule names the group or the logins it covers, the window — a number of days or a number of trades, opening at the client's first closed trade rather than at registration — whether it refunds the netted result or every losing trade with winners ignored, the percentage, the caps per client and in total, and the exclusions: trades held under a few seconds, bonus-funded accounts, accounts with an open abuse flag, and named instruments or logins.

Every figure is the gateway's, so the number an operator approves is the number the gateway attempts. There is no button that pays: **Preview** opens the plan — every client, the loss in their window, the arithmetic from that loss to a figure, and the caps it met — and the Pay button lives inside that dialog behind a typed confirmation naming the total. An exclusion register that cannot be read *holds* the client.

Rules · What it has paid · Journal · Kill switch, first on the page · A dash is a dash

Nothing is paid twice

Granting and applying a withdrawal's impact need the caller's own idempotency key, and the gateway will not invent one: a key it made up is one this console could not send twice, so every retry would be a second payment.

Converting is two legs, in order

Credit comes off first and balance goes on second. If the second leg fails the first is put back and the screen says so; if the restore also fails, it refuses to pretend and tells an operator that the account is short and has to be corrected by hand.

A grant can sit at **pending** — a real, stored state meaning the server was asked and never answered, shown as that and not as a failure.

12 HOW IT CONNECTS

One book, read once, by everything here

Nothing in this module keeps its own copy of the book, and nothing reads it twice. The exposure aggregation is computed in a single pass on the gateway and cached; the closed-trade history is normalised once, so a trade cannot be one duration on one screen and another on the next.

TAB	READS	WRITES
Dashboard	The shared exposure snapshot and the margin read	Its own limit rules; a breach is announced, never acted on
Device ID, IP Address, Arbitrage, Scalpers	The fraud engine's detector registry and case queue	Nothing; decisions are taken on Cases
Swap Abuser, Protection Rules	The same anti-abuse desk — flags, evidence and measures	Flag status, measures and their lifts
Dividend Abuser	One detector, and the Dividend Manager's ledger against position history	That detector's thresholds, in the shared settings document
The Anti Fraud tabs	Cases, detectors, scans and statistics	Settings, scan jobs, case decisions, measures
Exposure Limit	The same snapshot the Dashboard reads, including the group-by-instrument cross-tab	The limit set; a reaction closes positions
Anti Scalp	The closed-trade history, normalised as the fraud desk reads it	The policy; a reaction withdraws a trading permission
Auto-Hedger	The same snapshot with the coverage accounts excluded	Policy, and orders on the coverage account
Bonus & Credit, TradeShield	Campaigns, grants, closed volume and the exclusion registers	Credit and balance operations on a client's account

Permissions

Reading a desk and acting on it are separate grants. Closing positions, moving a group or cutting leverage each need the accounts permission on top of the desk's own, and anything that books money needs the balance permission. A refusal names both halves.

Licensing

Anti Fraud, Anti-Abuse Protection, the Auto-Hedger, the Risk Dashboard and the Bonus & Credit Manager are licensed components: a brokerage that has not bought one does not see its menu at all. Exposure Limit and Anti Scalp ride on the gateway itself.

Every desk here keeps its own journal, and every one records refusals beside successes. **"What did we do about that account, when, and who decided"** is a question this module answers by itself.