METHODOLOGY

How every number is made

A plain-language white paper: how the system works from start to finish, where the data comes from, how each metric is calculated, and what the models leave out.

Version 1.3 · dated 23 Nov 2026 · covers site version 91.2 (sections 1 to 9 as in v1.1)
Status: **provisional** under 300 counted trades. Results are in R and %; no dollar amounts are shown.
Changes: **v1.3 (23 Nov 2026)**: corrects the name, Mont for Monte Carlo and Ex for expectancy (v1.2 said Mon and Tex), and adds a short definition of a 1:1 trade under it. It changes no rule and no number.
v1.2 (23 Nov 2026): adds an opening overview, How MARS works, start to finish, with what the name means (Monte Carlo + expectancy, alpha, rail): the plan is sealed first, the Monte Carlo ruler is built and frozen before the first counted trade, and then each week the live record is placed against it. Sections 1 to 9 are unchanged. It changes no rule and no number.
v1.1 (20 Nov 2026): brought up to date with the site and the tracking workbook before week 1: risk per trade on the rail runs from 1% to 5.75% (6% only by a manual tier), where v1.0 said 1% to 5.5%; returns, drawdown and acceleration are flow-neutral, so a withdrawal is never counted as a loss; the Monte Carlo ruler is described as the workbook runs it (50,000 frozen paths, where v1.0 said 1,000); the health check, verification chain, R | % switch, manual-tier column, Alpha Cup handicap, Filters list and what is not modelled are brought in line with the code. It changes no rule and no number. Two sentences of the v1.0 page (win rate and RAPF wording) were updated on the site after it was sealed; the sealed v1.0 PDF and page source below are unchanged.
⬇ Download this version as PDF · Sealed 23 Nov 2026: SHA-256 seal and its proof · Version 1.2 (sealed 23 Nov): PDF, seal and its proof · Version 1.1 (sealed 20 Nov): PDF, seal and its proof · Version 1.0 (sealed 13 Nov): PDF, seal and its OpenTimestamps proof

Contents

How MARS works, start to finish

1. The data path
2. Units: R, % and the breakeven band
3. The five core metrics
4. The Monte Carlo ruler
5. Confidence and sample size
6. Verification

7. The Alpha Cup scoring
8. The Filters rules
9. What is not modelled

How MARS works, start to finish

MARS tests one question in public: does a written 1:1 trading plan have a real edge, and can a risk rail turn that edge into growth without a drawdown the account cannot afford? Everything happens in a fixed order, so the rules and the yardstick exist before the results do.

The name. MARS stands for Montex Alpha Rail System, and the name lists the method. *Mont* is Monte Carlo: the ruler, built and frozen before the first trade (step 2). *Ex* is expectancy: EV, the master metric that grades the live record (step 5). *Alpha* is beating that ruler without a worse drawdown, efficiency or compliance. The *rail* is what makes that possible: gates, tiers, day colour, weekly status and System Lock set every trade's risk (step 3). The ruler says what should happen; expectancy says what is happening.

What 1:1 means. Every trade is 1:1: the stop and the profit target sit the same distance from the entry, so a win is +1R and a loss −1R before costs, where R is the amount risked. With the 0.05R cost a trade, the break-even win rate is 52.5%; the plan assumes 60% until the record measures it. The site's Glossary entry *1:1* explains it from the start, in pips, for readers new to trading.

1. **Write the plan and seal it.** The Trading Plan and Risk Charter sets the rules before the first counted trade: the 1:1 strategy and its 14-item entry checklist, four trades a day on four trading days a week for 78 weeks from Mon 7 Dec 2026, and the risk rail that sets every trade's risk. Each version is fingerprinted and timestamped, so it cannot be quietly rewritten to fit the results.
2. **Build the ruler and freeze it.** Before go-live the plan itself is simulated 50,000 times: 60% wins, 0.05R of cost a trade and the rail's own sizing. The spread of those paths shows what a normal result looks like for every week of the test if the plan's assumptions are right. It is run once and never rerun (section 4).
3. **Trade on the rail.** Entries are the trader's judgement, checked against the 14 items; risk is not. Each evening the rail sets the next day's risk per trade from the drawdown gate, the tier, the day colour and the weekly status: 1% to 5.75% a trade, or a manual tier the record shows, and zero at System Lock.
4. **Record every trade the same day.** The broker statement feeds the tracking workbook, an independent health check and this site, and the new trades are sealed (sections 1 and 6).

Every result is net of costs, in R and % (section 2).

5. **Compare with the ruler.** Each week the real return and the deepest drawdown are placed against the frozen paths for that week, and the five core metrics are judged in order of authority: EV, drawdown, RAER, RAPF and equity acceleration (sections 3 and 4). Home › Pace and finish line shows each week inside the Monte Carlo cone. Beating the ruler without a worse drawdown, efficiency or compliance is alpha; beating it with more risk is not.
6. **Let the evidence build.** No rule, mode or risk setting changes before 30 counted trades. Milestone reviews at 30, 100 and 300 trades ask how strong the evidence for an edge has become (section 5), and every verdict stays provisional until 300. Any later change is dated, gets a new version of the Charter and is measured on Records › Rule-change impact.

1. The data path

1. **Broker statement.** After the last trade of the day, the account's detailed statement is saved from the trading platform. It is the single source of every trade.
2. **Tracking workbook.** The statement is loaded into the MARS tracking workbook, which matches each closed trade to its day and slot, and computes its net result in R after spread, commission, swap and slippage.
3. **Health check.** An independent script re-reads the statement and checks the workbook: every trade present once, net results to the cent, equity arithmetic, next-day risk. A failed check stops the end-of-day routine until it is fixed.
4. **Public page.** A build script turns the workbook into this site's data. Before anything is published, a self-test of the page's risk-rail maths, a glossary check and a script check must pass, and a leak check refuses any text that would show an account number, a balance or a dollar amount (the fixed R teaching example of 10 and 5 dollars is the only exception). Any failure keeps the last good version of the site.
5. **Seal.** Newly published trades are fingerprinted and stamped in the verification chain (section 6).

2. Units: R, % and the breakeven band

- **R** is the amount risked on a trade. A trade's net R = its net profit or loss ÷ the amount it risked. Because every trade is 1:1, a clean win is about +1R and a clean loss about −1R, less costs.

- **%** is the change in account equity. Equity results depend on the risk the rail set that day (1% to 5.75% a trade; 6% only by a manual tier, see the Charter), so the same R can mean different %. Returns are **flow-neutral**: each day's own return, chained from the start, so money withdrawn or deposited is never counted as a gain or a loss. Withdrawals are shown separately, in %, on Records › Statements.
- **R | % switch.** On charts that carry it (view level 3 and up), % = each trade's R × that day's risk per trade × the balance at the start of that day, in points of the starting balance, so the figures add up to the account's growth. Edge, probability and benchmark charts have no switch.
- **Breakeven band.** A result between −0.10R and +0.10R is a breakeven, not a win or a loss. Win rate = wins ÷ (wins + losses); breakevens are left out of it but count in every R total. Equilibrium itself reports wins ÷ all trades, which counts breakevens as not won; Home shows that figure under the Win rate tile, and the replication kit gives both. The R profit factor (R won ÷ R lost) also leaves breakevens out.
- **Costs** are always inside the numbers: every R and % on the site is net of all costs.
- **Overrides.** The Daily record has an Override column: a dash when the rail set the day's tier, or the manual tier used against the rail's tier (for example "Manual T4 ↓ (rail T5)"), marked when it was raised before 30 counted trades.

3. The five core metrics, in order of authority

Metric	Question	Formula
EV (expectancy)	Is there an edge?	Average net R per trade.
Drawdown	Can the account survive it?	Fall from the Equity Peak High (the highest day close so far) to the day's close, in %, on flow-neutral equity. The gate reads today's drawdown; the core-metric grade reads the deepest so far.
RAER	Is the risk put on turned into return?	Total profit ÷ total risk put on, both in account money: each trade counts at the risk and account size it was taken at.
RAPF	Is the profit quality real?	Profit factor (R won ÷ R lost, breakevens within ±0.10R left out) × a 3.5% reference risk ÷ the average risk actually used: profit quality at the plan's reference size. Whether bigger positions won more is RAER's question, not RAPF's.
Equity acceleration	Is compounding speeding up?	The last traded week's flow-neutral % return minus the week before.

The order matters: a fast-growing account with no edge or with a drawdown it cannot survive fails the test, whatever its speed.

4. The Monte Carlo ruler

The benchmark is a model of the plan, not a target. Its main ruler is a frozen simulation of 50,000 paths of the plan itself: 1:1 trades at a 60% win rate with 0.05R of cost on each, four trades a day and four days a week for 78 weeks, with risk set day by day by the rail's rules for the mode in use (gates, tiers, day colour, weekly status, the caution lock and Balanced's T6 after an Above-plan week). It was run once before go-live with a fixed seed and is never rerun, so the ruler cannot move. Each week the live equity and its deepest drawdown are placed against the 25th, 50th and 75th percentiles of those paths for that week.

EV and RAER are judged against the plan's EV of +0.15R a trade, with a 25th-to-75th band of ±0.6745 standard errors for the number of trades; RAPF against the plan's profit factor over the same win-rate band. A second, older benchmark (the 7-Tier Monte Carlo, 50,000 paths of MARS's original payoffs, which size below Balanced) is shown beside it as a check, not the grade. Statuses start at 30 trades.

Alpha means beating the benchmark without a worse drawdown, worse efficiency or worse compliance. Beating it with more risk is not alpha.

5. Confidence and sample size

- EV is shown with a 95% range: EV ± t × standard deviation ÷ √trades. The site's edge cards use the t value, about 2 and a little wider with few trades; the workbook's headline range and the slice breakdowns use 1.96. While the range includes zero, the edge is not proven.
- Win rates use the Wilson 95% range, which behaves well with few trades.
- Statuses wait for 30 trades; a firm verdict waits for 300. Until then every result carries a provisional label.

6. Verification

Each update that publishes new trades adds a numbered block file holding the SHA-256 fingerprint of those trade rows, the fingerprint of the broker statement at that moment (the statement itself stays private) and the fingerprint of the previous block, so the blocks form a chain. The blocks are stamped with OpenTimestamps, which anchors them in the Bitcoin blockchain. The Integrity page shows, for each trading day, the block that fixed it, its Bitcoin time and whether it was provably fixed before the next trading day's first trade (allowing two hours for Bitcoin's loose clock). Anyone can download a block and its proof and check it at opentimestamps.org. A replication kit (trades as CSV, and an Excel workbook with live formulas) lets a reader rebuild the headline numbers.

7. The Alpha Cup scoring

The Alpha Cup lets a visitor's own settings play the live record on the same trading days. A match is scored on 19 statistics in five groups, weighted in the MARS order of authority:

Group	Statistics	Points each
Edge	EV per trade, cost cushion, SQN, chance the edge is real	5
Survival	Max drawdown, longest losing run, time under water, ulcer index	4
Efficiency	RAER, return ÷ drawdown, smoothness	4
Quality	Risk-weighted profit factor, R profit factor, payoff ratio, green days	2
Speed	Return, equity acceleration, new equity highs, best day	1

- Each statistic goes to the side with the better value (higher or lower, by statistic); equal shown values split the points. A statistic scores only once both sides have its minimum number of trades. The Cup's quality statistic weights each trade by its risk; MARS's own RAPF is on Edge › RAER & RAPF.
- Bonus plays: +3 for each finished week with trades and no fall through a gate (none in or after a System Lock week); −10 for reaching System Lock.
- Handicap: a team whose edge before costs is above +0.25R a trade is optimistic and gives the record 5 points; above +0.40R it is fantasy and gives 15.
- The visitor's trades are drawn at random from their own settings; the live record's trades are the real ones and never change. The Scouting report replays the team with fresh luck for 300 or 1,000 seasons to show how much of a result is luck.
- Prize weeks count only with the team registered before the week began, so the past cannot be rewritten. Everything stays in the visitor's browser. Illustration only.

8. The Filters rules

The Filters panel narrows 36 trade-level pages (for example edge, breakdowns, excursions, execution and the trade log) to the trades that match a chosen asset class, market, session, timeframe, direction, weekday or slot, and an optional date window; the market list follows the chosen asset class. It never touches pages built from the whole record in sequence, and those pages also ignore the date window: account equity, drawdown, the Monte Carlo benchmark, gates and tiers, simulations, records and statements, and the Alpha Cup. A filtered view is a slice, not a different track record: under 30 matching trades the panel says the slice is mostly luck.

9. What is not modelled

- The site's own simulations and replays use the Balanced rail with its gates, tiers, day colour, weekly status and T6 after an Above-plan week, and stop on the trade that reaches System Lock. The caution lock is modelled only where a risk mode is chosen (Mode comparison and the Lab's mode lab, tournament, theatre and console) and in the frozen benchmark. Manually chosen tiers and holiday notes are modelled nowhere.
- The frozen benchmark checks the System Lock at each day's close, assumes four trades on every trading day, and stops a path once it reaches the account's goal.
- Simulated trades are independent coin flips at the chosen win rate; real trades can cluster.
- Costs in simulations are a fixed average; real costs vary by pair, session and stop size.
- Past order of trades only: what-if tools replay the record as it happened and cannot know how other choices would have changed later trades.

> Education and record-keeping only; not investment advice. See Disclaimer and privacy and the Trading Plan and Risk Charter.

To check the seal yourself: the page you are reading is wrapped in the site menu, so hash the exact sealed files instead: the PDF and the sealed page source (v1.3 page source · v1.2 page source · v1.1 page source · v1.0 page source). Their SHA-256 must match the seal file.