

Quarterly Config Watch #1

October 9, 2026 · MarketMania Research · Quarterly series

QUARTERLY · CONFIG WATCH · Issue #1 · one calendar quarter, 60 replays, one survival curve · Issue date October 9, 2026 · Language English · Methodology v1.1 + Config Watch amendments (approved 12.08, hash e66c7e8c864a2233)

Search window: **QUARTER** Jul 1-Sep 30, 2026 -- the frozen calendar quarter (Wed Jul 1 00:00 -> Wed Sep 30 23:59 UTC, 92 days; platform data from Wed Jul 15 18:01 UTC, 78 days), **Q3 2026 (partial: platform data from Jul 15)**; cut-off edge Thu Oct 1 00:00 UTC, so slots dated Oct 1 fall outside it; report cut-off Thu Oct 1, 2026, 16:00 UTC. Slots fall inside the window; trades settle past its right edge at their horizon -- the same 2-3 day luft the engine has used since Config Watch #1 and that the live presets run on.

Prior issue: none -- this is the first quarterly issue of the series. The two monthly issues inside the quarter are Monthly Config Watch #1 (September 4, 2026, calendar August, Aug 1-31, 2026: best mature council +\$1,618.4, best mature solo +\$1,721.9, 20 of 26 published cards net-positive) and Monthly Config Watch #2 (October 3, 2026, calendar September, Sep 1-30, 2026: +\$1,041.5, +\$923.2, 17 of 53), and every number below that sets the quarter beside a month is read from their open data. Config Watch #9 carries the calendar week Sep 28-Oct 5, 2026 and the four weeks Sep 7-Oct 5, 2026 from two collects of its own; Monthly Config Watch #2 the calendar September from a collect of its own; this issue's collect sweeps the calendar quarter once, and no weekly or monthly board is reprinted here.

RESEARCH SNAPSHOT

10,429 simulations, 0 errors, in one frozen collect	QUARTER Jul 1-Sep 30, 2026 (data from Jul 15)	cut-off edge Thu Oct 1 00:00 UTC ; report cut-off Thu Oct 1, 2026, 16:00 UTC	7-stage search: scan -> grid -> refine -> stake sweep -> per-ticker -> OOS replays
engine 1.1; \$1,000 deposit, 10x leverage	stake is universe-aware from sim #1 (table below)	maturity gate: >=30 closed (QUARTER)	criteria frozen in writing before launch (w309 plan)
walk-forward: 56 past cards + 4 live presets replayed verbatim on the quarter	30 of 56 past cards net-positive on the quarter	stage-2 grid: 9,028 cells on this window	models: 7 (grok axis = grok-4.6 (incl. 4.5-era))
best QUARTER council +\$2,271.8 (+227.2%)	best QUARTER solo +\$1,491.0 (+149.1%)	zero-trade cells: 951 (whole collect)	margin-skip alerts on leaders: 14 rows (QUARTER)

KEY FINDING · OBSERVATION (one calendar quarter)

The best mature council on the calendar quarter (Q3 2026) makes +\$2,271.8 (+227.2% of a \$1,000 deposit, 63.9% win rate, 74.3% maxDD, 427 closed and 23 unfunded entries), 30 of the 56 cards this series has published are net-positive on the same quarter, and 9 of the 17 cards that survived both fresh windows of Config Watch #9 are among them.

MONTHLY EVIDENCE -> QUARTERLY VERDICT (rule below)

Metric	M1 Jul 15-31	M2 Aug 1-31	M3 Sep 1-30	Quarter	Verdict
Published cards net-positive (majority)	n/a	20/26	17/53	30/56	Inconclusive (1/2)
Best council universe	n/a	SOL+XRP	BNB+XRP	SOL+XRP	Inconclusive (1/2)
Calibration helps (matched pairs)	n/a	37.9%	65.6%	56.1%	Inconclusive (1/2)


Verdict rule: each row states one reading of the quarter; a month agrees when its own monthly issue shows the same reading on that calendar month (same side of the line, or the same name). Confirmed = 2 of the 2 comparable months agree with the quarter; Inconclusive = 1; Rejected = 0 -- the monthly ladder (Confirmed = 3 or 4 of 4; Inconclusive = 2; Rejected = 0 or 1) brought to two comparable columns. A cell that reads n/a has no comparable number in that month: it never agrees and leaves the count. M1 (Jul 15-31) has no monthly Config Watch -- the series' first issue is dated Aug 13 -- so it reads n/a in every row and every row is scored out of two. Month cells come from the two published monthly issues' own open data (M2 = Monthly Config Watch #1, M3 = Monthly Config Watch #2), each on its calendar month; the quarter column is this collect's calendar-quarter pull.

Sources, cell by cell. Row 1 counts the replays each monthly issue ran on its own calendar month -- Monthly Config Watch #1 20 of 26 (Aug 1-31, 2026), #2 17 of 53 (Sep 1-30, 2026) -- against 30 of 56 on the quarter in the walk-forward table below. Row 2 reads the coin universe of each monthly issue's council card #1, and here the top council card of the quarter. Row 3 is the share of matched calibrate-ON / calibrate-OFF pairs the axis improved on that month, over the whole stage-2 grid in every column: 1,419 pairs in Monthly #1 and 1,446 in #2, against 3,497 whole-grid pairs on the quarter. 3 cells read n/a, all in M1: July has no monthly Config Watch issue. The two months cover 61 of the quarter's 78 data days; Jul 15-31 lie in no monthly issue's window.

TL;DR

- OBSERVATION -- 30 of the 56 published cards are net-positive on the calendar quarter** (calendar August: 20 of 26 in Monthly Config Watch #1; calendar September: 17 of 53 in Monthly Config Watch #2). Config Watch #9 read the same 56 cards on its fresh week (Sep 28-Oct 5, 2026): 25 positive; on its four-week window (Sep 7-Oct 5, 2026): 27. By generation on the quarter: Monthly-2 1/3, issue-8 1/5, issue-7 3/5, issue-6 2/5, issue-5 1/4, issue-4 3/5, Monthly-1 3/3, issue-3 3/6, issue-2 7/8, issue-1 6/12; live presets 4/4. None of the 56 cards was published before the quarter began (issue-1, the first, on Aug 13). The quarter holds 3 of that week's 7 days and 24 of that four-week window's 28, and it contains the search window of every cohort, so the readings are not independent tests.
- The quarter's ceiling is +\$2,271.8 on the council board and +\$1,491.0 on the solo board** (calendar August: +\$1,618.4 and +\$1,721.9; calendar September: +\$1,041.5 and +\$923.2). The council card runs at +227.2% of a \$1,000 deposit, 63.9% win rate and 74.3% maxDD on 427 closed trades with 23 unfunded entries; the solo card at +149.1%, 60.9% win rate and 80.8% maxDD on 443 closed with 46 unfunded entries. The grid behind them: median -\$195.1 and 25.5% positive across 9,028 stage-2 cells, 491 of them above +\$500, range -\$908.7 to +\$2,271.8. The two monthly grids ran 41.8% (August) and 29.6% (September) positive; Monthly Config Watch #1's card #1 reads +\$1,000.1 on the quarter against +\$1,618.4 printed on its month and Monthly Config Watch #2's card #1 reads -\$619.4 on the quarter against +\$1,041.5 printed on its month (section below).
- The survival curve reaches +8 fresh weeks for the issue-1 generation** (+7 in Monthly Config Watch #2). Share of a generation's cards net-positive on its latest fresh week: issue-1 8.3% (+8); issue-2 25.0% (+7); issue-3 33.3% (+6); issue-4 60.0% (+5); Monthly-1 66.7% (+5); issue-5 75.0% (+4); issue-6 60.0% (+3); issue-7 80.0% (+2); issue-8 40.0% (+1); Monthly-2 100.0% (+1); every generation has at least one fresh week. Every cell is read out of the issue that ran that replay, and the points sit on 8 different calendar weeks, from Aug 10-17, 2026 to Sep 28-Oct 5, 2026.
- The 17 survivors and the 39 faded cards of Config Watch #9, side by side.** The cards that cleared both of its fresh windows against those that did not: median quarter maxDD 80.2% vs 87.0%, median win rate 63.9% vs 61.8%, 4 of 17 against 15 of 39 on the 1d horizon, and 5 of 17 (29.4%) vs 5 of 39 (12.8%) calibrated. 17 rows against 39 rows on two windows is a description of two groups, not a test, and none of the traits is claimed to cause survival.

Survival by generation -- cards net-positive on the k-th fresh week



Share of each generation's cards net-positive on the k-th fresh week after its own window edge (issue-1 12 cards, issue-2 8 cards, issue-3 6 cards, issue-4 5 cards, Monthly-1 3 cards, issue-5 4 cards, issue-6 5 cards, issue-7 5 cards, issue-8 5 cards, Monthly-2 3 cards). Each point is one published walk-forward table: issue-1's eight points are Config Watch #2 to #9, Monthly-2's one point is #9; the count behind every point is in the survival table below. The x axis counts fresh weeks, not calendar weeks -- the points sit on 8 different weeks, from Aug 10-17, 2026 to Sep 28-Oct 5, 2026 -- and neither the live presets (no publication date) nor the quarter column (a 92-day window) is plotted.

Why it matters

MarketMania publishes default LLM-council trading configs. The weekly Config Watch asks what one week of large-scale in-sample search extracts and how much of it survives the next week, and the monthly issues ask it of the calendar month. This issue asks the same question with the calendar quarter as the unit: one frozen 92-day window (78 days with data), the same engine, the same criteria and the same frozen search plan, and the same walk-forward table read on that window. Two measurements the monthly issues carry are standing sections here too: the survival curve by generation, which places each published card's replay on the k-th fresh week after its own issue and now covers 10 generations, and the structural comparison of the 17 cards that cleared both fresh windows of Config Watch #9 against the 39 that did not. The quarter is not an independent test of the months or the weeks: 61 of its 78 data days lie in the two monthly issues' months, it contains the search window of every published card, and every section that reads them together says so.

How we searched

One frozen collect, run to the plan with the criteria frozen in writing *before* launch (w309 plan) and printed into the meta of the JSON deliverable: the quarter run, which sweeps the calendar quarter in one pass. Every stage count below is that run whole. It is a run of its own -- no part of it ships with a weekly or a monthly issue -- and the stage-2 grid is 9,028 cells on its one window.

- **1. Scan (s1)** -- 476 sims across council compositions and coarse settings.
- **2. Systematic grid (s2)** -- 9,028 sims sweeping the declared axes (archetype, FH/TF, entry filters, universe, council size, membership, TP/SL source, ladder, break-even, TP/SL shift, calibrate and stake) -- the denominator for every distribution claim below.
- **3. Refinement: ladder + stake sweep (s2b+s3b)** -- 123 sims around the grid winners.
- **4. Preset neighbourhood grid (s2p)** -- 53 sims: one-knob neighbours of the live presets that have a local grid this issue.
- **5. Per-ticker probes (s5t)** -- 60 sims: best finalists split onto single coins.
- **6. Dedicated per-ticker (s6t+s6tc)** -- 585 sims: a compacted independent per-coin grid plus calibrate twins (the standing branch from Config Watch #2).
- **7. OOS replays (s4cwm2+s4cwm2o+s4cw8+s4cw8o+s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)** -- 104 sims: every published card and all four live presets, verbatim, on the one window of this collect, plus origin re-runs of the cards of Monthly #2, issue #8, issue #7, issue #6, issue #5, issue #4, Monthly #1, issue #3 and issue #2 for the drift baseline.

Total 10,429 simulations, 0 errors, 951 zero-trade cells; stage sums reconcile exactly inside the collect. The stage-2 grid is 9,028 cells on this window against 4,112 on Monthly Config Watch #1's calendar month and 4,176 on Monthly Config Watch #2's. Equity series, the grid histogram, the baselines and the calibrate pack are produced inside the same collect -- there is no separate extras pass.

Stake: universe-aware from simulation #1 (\$1,000 deposit)

Instruments in universe	Stake cap	Sweep values also simulated
1 coin	\$450	\$300 / \$375
2 coins	\$450	\$300 / \$375
3 coins	\$300	\$225 / \$375
4 coins	\$250	\$200 / \$300
5 coins	\$200	\$150 / \$250

A \$1,000 deposit cannot fund five concurrent \$450 positions. Stake is a function of how many instruments the config trades, recomputed on every universe change and swept as its own axis; any simulated stake above the cap for its universe size is marked **concentrated** and excluded from the headline boards (it stays in the open data). Every card also prints **skipped_nofunds** -- entries the engine could not fund. Margin-skip alert rows on this window's leaders: 14, every alert the collect raised on its one window; council card #1 (23 unfunded entries) and solo card #1 (46 unfunded entries) are among them, so their economics are not as printed (methodology SS2).

Models under test (exact engine IDs): claude-fable-5, claude-opus-5, gpt-5.6-sol, deepseek-v4-pro, gemini-3.1-pro, qwen-3.8-max, grok-4.6. The grok axis runs as **grok-4.6 (incl. 4.5-era)** -- one lineage; the grok-4.5 -> 4.6 cutover dated Aug 24 in the search plan

falls inside the quarter (Jul 1-Sep 30, 2026), and published grok-4.5 configs replayed out-of-sample get the same 4.5 -> 4.6 lineage remap ([remap] rows). The lineage probe of the collect counts 10 raw model IDs with forecasts in the quarter -- the 7 current ones and 3 older ones, each the predecessor of a current ID: claude-opus-4.8 (6,315 rows, last slot Jul 30 08:01) -> claude-opus-5 (first slot Jul 30 08:01); qwen-3.7-max (7,479 rows, last slot Aug 5 09:01) -> qwen-3.8-max (first slot Aug 5 10:01); grok-4.5 (11,128 rows, last slot Aug 24 09:01) -> grok-4.6 (first slot Aug 24 10:01).

TOP councils -- calendar quarter (Jul 1-Sep 30, 2026)

#	Council	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	fable, opus, gpt, gemini	SOL+XRP	4h/1h	+\$2,271.8	+227.2%	74.3%	427 *	63.9%
#2	fable, opus, deepseek, gemini	BNB+XRP	1d/1d,4h	+\$1,107.1	+110.7%	17.8%	72	93.1%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	>=60 (each)	farthest	3/1	reenter	\$450
#2	50:40,100:60	on	TP -10 / SL +75	>=50 (each)	median	3/1	reenter	\$450

Against the 9,028-cell stage-2 grid on this window: card #1 is the same simulation as the raw net-board leader (i=1379) and the top cell of the grid; card #2 lands in the +\$1,000 to +\$1,250 bucket (69 cells) with 98.9% of the grid finished below it. Grid median -\$195.1, 491 of the 9,028 cells above +\$500. Net % = net PnL as a percentage of the fixed \$1,000 starting deposit; every simulation runs at **10x leverage** (liquidation modeled by the engine). Entry = max_sideways / max_diff_side.


Two councils are published on this window. All ten slots of the mature net board are the same 2-coin universe (SOL+XRP), so the pairwise-different-universes rule admits exactly one; #2 is the best mature multi-coin council outside that universe on the quality boards -- it comes off the **Calmar** board (+\$1,107.1 on BNB+XRP, 17.8% maxDD, R2 0.932, no unfunded entry). Read the cards as two universes deep, not three. Card #1 is a configuration Config Watch #7 published: cw7-wk1, query for query, window aside, and its walk-forward row below reads the same numbers on the quarter.

The two cards answer the calibrate axis differently: card #1 is an uncalibrated cell and card #2 a **calibrate=ON** cell (Calibrate axis below). Card #1 runs the \$450 two-coin table stake and card #2 the \$450 two-coin stake; card #1 carries 23 unfunded entries and card #2 no unfunded entry, and card #1 draws down 74.3% on 427 closed trades against card #2's 17.8% on 72. The mature pool on this window is 7,126 configs at the >=30-trade gate (78 data days), against 3,265 on calendar August and 2,967 on calendar September at the monthly >=20 gate.

OPEN IN SANDBOX marketmania.ai/s/cwq1-q1 marketmania.ai/s/cwq1-q2


Each link opens this exact frozen window; results visible without sign-in (embedded share signature). The codes are minted by the publish script before this file goes live; every card's full frozen query signature is in the open-data JSON next to its simulation index.

Published card equity -- QUARTER (Jul 1-Sep 30, 2026) -- \$1,000 at 10x -- 427 settled trades



Daily settled equity for card #1 above (\$1,000 start, 10x leverage, frozen window; #1 red -- fixed rank colors across all issues). A curve running past the window end is an open-at-cutoff trade settling at its horizon (the 2-3 day left); the last point reconciles to the card's net to the cent. There is no muted contrast line on this window: the mature leader is also the net-board maximum (the same simulation, i=1379), so a second curve would be the same line. The line peaks at +\$2,688.0 on 2026-09-22 and ends the window at +\$2,271.8. Card #2, the solo card and the other board leaders keep their daily series in the open data -- the PDF carries one equity chart.

Stage-2 systematic grid -- net PnL distribution -- QUARTER window (Jul 1-Sep 30) -- n=9,028



Net PnL across all 9,028 stage-2 grid sims, QUARTER window (Jul 1-Sep 30, 2026), universe-aware stake throughout. Grid median **-\$195.1**, 25.5% positive (calendar August: \$0.0, 41.8% positive; calendar September: -\$5.7, 29.6% positive); Config Watch #9's grids ran 30.3% positive on its fresh week and 36.3% on its four-week window. The \$0 bin holds 2,197 cells; 950 sims on this window produced no trade at all. Range -\$908.7 to +\$2,271.8; 491 cells finished above +\$500 and the maximum of the grid is council card #1. Red marker = best solo default at table stake; green = the published cards.

The monthly cards against the quarter

Monthly Config Watch #1 published two councils and one solo on calendar August (Aug 1-31, 2026, 31 days) and Monthly Config Watch #2 published two councils and one solo on calendar September (Sep 1-30, 2026, 30 days); this issue publishes two councils and one solo on the calendar quarter (Jul 1-Sep 30, 2026, data from Jul 15). Both months lie inside the quarter -- 61 of its 78 data days. The table sets each monthly card's own month, as its issue published it and as this collect re-runs the same query on the same month today, beside the card's replay on the quarter; this issue's three cards have no month of their own and read n/a there.

Card	Council / solo	Coins	FH/TF	Published	Re-run	Quarter	maxDD Mo/Q	Cls Mo/Q
cwm1-m1	fable, gemini	SOL+XRP	1d/1d	+\$1,618.4	+\$1,618.4	+\$1,000.1	40.5% / 153.9%	31 / 79
cwm1-m2	fable, gpt, deepseek, gemini, qwen	ETH+ SOL+ BNB+XRP	1d/1d	+\$756.6	+\$451.2	+\$91.5	8.1% / 53.9%	41 / 116
cwm1-s1	gemini	SOL+XRP	1d/1d,4h	+\$1,721.9	+\$1,762.0	+\$856.8	54.9% / 185.8%	43 / 104
cwm2-m1	fable, opus, gemini, qwen	BNB+XRP	4h/1h	+\$1,041.5	+\$1,041.5	-\$619.4	37.7% / 95.2%	97 / 112
cwm2-m2	fable, opus, gemini, qwen	SOL+XRP	4h/1h	+\$1,021.7	+\$1,021.7	+\$1,047.4	47.7% / 142.1%	90 / 239
cwm2-s1	gemini	SOL+XRP	4h/1h	+\$923.2	+\$487.3	-\$575.4	32.0% / 87.7%	207 / 140
cwq1-q1	fable, opus, gpt, gemini	SOL+XRP	4h/1h	n/a	n/a	+\$2,271.8	n/a / 74.3%	n/a / 427
cwq1-q2	fable, opus, deepseek, gemini	BNB+XRP	1d/1d,4h	n/a	n/a	+\$1,107.1	n/a / 17.8%	n/a / 72
cwq1-s1	gemini	SOL+XRP	4h/1h	n/a	n/a	+\$1,491.0	n/a / 80.8%	n/a / 443

Published = the net the card's own issue printed for its month; Re-run = the same query on the same month in this collect; Quarter = this collect's replay of the published query (for this issue's cards, the card itself). Paired columns read **the re-run month first, the quarter second**; CIs = closed trades. **cwm1-s1 is cw3-s2, cwm2-m1 is cw8-m302, cwm2-m2 is cw8-m301 and cwq1-q1 is cw7-wk1**, query for query, window aside -- their replay rows below carry the same quarter numbers; the nine rows are nine distinct configurations. Place on the quarter's mature net board (each branch's top 10, as the collect ranked them): cwq1-q1 1st of the councils; cwq1-s1 1st of the solos; cwm1-m1, cwm1-m2, cwm1-s1, cwm2-m1, cwm2-m2, cwm2-s1 and cwq1-q2 are not on it. cwq1-q1, cwq1-q2 and cwq1-s1 have no month of their own, so Published and Re-run read n/a. Boards, the replay set and the origin re-runs are what the collect exports, not every cell of the grid on every month.


All six monthly cards can be read on their own month and on the quarter, and 5 of them make less on the quarter than on their own month. cwm1-m1 makes +\$1,618.4 on calendar August and +\$1,000.1 on the quarter (-\$618.3, on 79 closed trades against 31); cwm1-m2 makes +\$451.2 on calendar August and +\$91.5 on the quarter (-\$359.6, on 116 closed trades against 41); cwm1-s1 makes +\$1,762.0 on calendar August and +\$856.8 on the quarter (-\$905.2, on 104 closed trades against 43); cwm2-m1 makes +\$1,041.5 on calendar September and -\$619.4 on the quarter (-\$1,661.0, on 112 closed trades against 97); cwm2-m2 makes +\$1,021.7 on calendar September and +\$1,047.4 on the quarter (+\$25.6, on 239 closed trades against 90); cwm2-s1 makes +\$487.3 on calendar September and -\$575.4 on the quarter (-\$1,062.6, on 140 closed trades against 207). The gaps run from \$25.6 to \$1,661.0. The re-run reproduces the published figure for cwm1-m1, cwm2-m1 and cwm2-m2; cwm1-m2 reads +\$451.2 against the +\$756.6 its issue printed, cwm1-s1 reads +\$1,762.0 against the +\$1,721.9 its issue printed and cwm2-s1 reads +\$487.3 against the +\$923.2 its issue printed (drift, walk-forward section).

Walk-forward on the calendar quarter

Every config this series published up to Monthly Config Watch #2 -- all 56 cards from Config Watch #1 to #8 and Monthly Config Watch #1 and #2 -- replayed VERBATIM (same knobs, same stake as printed, no re-tuning) on this issue's calendar quarter, plus the four hand-tuned presets running live on the platform, listed first. The replays and their origin re-runs cost 104 simulations inside this collect, which replays every row on its one window. It is the board Config Watch #9 read on its two fresh windows; Config Watch #9's own five linked cards (published Oct 7) are not on it: their search windows (Sep 28-Oct 5, 2026 and Sep 7-Oct 5, 2026) end after the quarter does, and this collect's plan carries no replay stage for them.

Verdict rule, on the one window this issue has. Survived = net-positive on the calendar quarter; **Faded** = net-negative. It is the standing rule (net-positive on every fresh window) with one window in it, and it is therefore easier to pass than the two-window version Config Watch #2 to #9 applied. The column is not out-of-sample: the quarter contains the search windows of all ten cohorts (days inside the quarter's data: Monthly-2's 30-day window 30; issue-8's 28-day window 28 and 7-day window 7; issue-7's 28-day window 28 and 7-day window 7; issue-6's 28-day window 28 and 7-day window 7; issue-5's 28-day window 28 and 7-day window 7; issue-4's 28-day window 28 and 7-day window 7; Monthly-1's 31-day window 31; issue-3's 28-day window 28 and 7-day window 7; issue-2's 30-day window 30 and 7-day window 7; issue-1's 30-day window 28 and 7-day window 7), so for all 56 rows the number a card shows here is partly scored on the days it was selected on. Read the column as decay context over a 92-day window with 78 days of data.

Config Watch #8's five published cards: home number vs calendar-quarter replay



Config Watch #8's five published cards: the number its own issue printed (grey, in-sample on its own window) against the same configuration replayed verbatim on this calendar quarter (green). 1 of the five is net-positive on the quarter. The quarter holds 28 of the 28 days of the m301 and m302 cards' window and 7 of the 7 days of the s1, wk1 and wk2 cards' window, so the green bars are scored partly on days the cards were selected on.

Config	Family	Coins	Home	New QUARTER (n)	Verdict
preset 4h4	fable, qwen, deepseek, opus	ETH+SOL	-- live	+\$987.9 (117)	Survived
preset bc2	gemini, gpt, deepseek	SOL+XRP	-- live	+\$1,460.6 (58)	Survived
preset sol	deepseek, fable, opus	SOL [conc]	-- live	+\$2,766.4 (110)	Survived
preset xrp	gemini, gpt, deepseek	XRP [conc]	-- live	+\$1,666.6 (27)	Survived

Config	Family	Coins	Home	New QUARTER (n)	Verdict
cwm2-m1	fable, opus, gemini, qwen	BNB+XRP	+\$1,041.5	-\$619.4 (112) +	Faded
cwm2-m2	fable, opus, gemini, qwen	SOL+XRP	+\$1,021.7	+\$1,047.4 (239) +	Survived
cwm2-s1	gemini	SOL+XRP	+\$923.2	-\$575.4 (140) +	Faded
cw8-m301	fable, opus, gemini, qwen	SOL+XRP	+\$1,018.6	+\$1,047.4 (239) +	Survived
cw8-m302	fable, opus, gemini, qwen	BNB+XRP	+\$943.0	-\$619.4 (112) +	Faded
cw8-s1	gemini	BNB+XRP	+\$445.4	-\$577.2 (46) +	Faded
cw8-wk1	gpt, gemini, qwen	SOL+XRP	+\$583.4	-\$602.5 (167) +	Faded
cw8-wk2	gpt, gemini, qwen, grok	BNB+XRP	+\$570.6	-\$566.3 (25) +	Faded
cw7-m301	fable, gpt, gemini, qwen	SOL+XRP	+\$1,351.8	+\$1,173.0 (364) +	Survived
cw7-m302	opus, gpt, gemini, qwen	BTC+ETH	+\$528.4	-\$572.6 (17) +	Faded
cw7-s1	gemini	SOL+XRP	+\$1,022.0	+\$1,084.5 (449) +	Survived
cw7-wk1	fable, opus, gpt, gemini	SOL+XRP	+\$1,467.8	+\$2,271.8 (427) +	Survived
cw7-wk2	fable, opus, gpt	BTC+ETH	+\$953.2	-\$592.5 (133) +	Faded
cw6-m301	opus, deepseek, gemini	SOL+XRP	+\$1,793.8	+\$1,841.7 (374) +	Survived
cw6-m302	gpt, deepseek, gemini, grok	BTC+ETH	+\$547.8	-\$376.8 (64) +	Faded
cw6-s1	grok	SOL+XRP	+\$404.8	-\$570.0 (103) +	Faded
cw6-wk1	opus, gpt, grok	SOL+XRP	+\$609.5	+\$1,325.8 (346) +	Survived
cw6-wk2	fable, opus, deepseek	BTC+ETH	+\$135.4	-\$609.5 (120) +	Faded
cw5-m301	fable, opus, gpt, gemini	SOL+XRP	+\$1,717.1	+\$1,133.7 (75) +	Survived
cw5-s1	gemini	BNB+XRP	+\$178.7	-\$64.1 (327) +	Faded
cw5-wk1	gpt, gemini	SOL+XRP	+\$510.4	-\$595.6 (162) +	Faded
cw5-wk2	fable, opus, gpt, gemini, grok	BTC+ETH	+\$183.5	-\$584.0 (90) +	Faded
cw4-m301	gpt, deepseek, gemini	SOL+XRP	+\$1,676.0	+\$1,651.2 (92) +	Survived
cw4-m302	fable, gpt, deepseek, gemini	BNB+XRP	+\$1,155.2	+\$1,034.8 (76) +	Survived
cw4-s1	deepseek	BTC+ETH+ SOL+BNB+XRP	+\$325.0	-\$166.3 (141) +	Faded
cw4-wk1	opus, deepseek, gemini, qwen	BTC+ETH+ SOL+BNB+XRP	+\$399.4	+\$418.4 (185)	Survived
cw4-wk2	fable, deepseek, qwen	BNB+XRP	+\$128.2	-\$570.9 (94) +	Faded
cwm1-m1	fable, gemini	SOL+XRP	+\$1,618.4	+\$1,000.1 (79) +	Survived
cwm1-m2	fable, gpt, deepseek, gemini, qwen	ETH+SOL+ BNB+XRP	+\$756.6	+\$91.5 (116) +	Survived
cwm1-s1	gemini	SOL+XRP	+\$1,721.9	+\$856.8 (104) +	Survived
cw3-m301	opus, deepseek, gemini	SOL+XRP	+\$2,137.6	+\$1,237.3 (279) +	Survived
cw3-m302	fable, opus, deepseek, gemini, qwen	BTC+ETH	+\$412.3	-\$370.0 (38) +	Faded
cw3-s1	gemini	SOL+XRP	+\$1,604.3	-\$570.8 (40) +	Faded
cw3-s2	gemini	SOL+XRP	+\$1,449.0	+\$856.8 (104) +	Survived
cw3-wk1	opus, deepseek, gemini	SOL+XRP	+\$2,175.1	+\$1,237.3 (279) +	Survived
cw3-wk2	opus, deepseek, gemini, grok	BTC+ETH	+\$837.6	-\$607.0 (90) +	Faded
cw2-30d1	opus, gpt, deepseek, gemini, qwen	BTC+ETH+ SOL+BNB+XRP [conc]	+\$744.7	+\$1,272.5 (122) +	Survived

Config	Family	Coins	Home	New QUARTER (n)	Verdict
cw2-30d2	opus, deepseek, gemini	ETH+SOL+BNB+XRP [conc]	+\$655.1	+\$1,292.9 (146) +	Survived
cw2-30d3	opus, deepseek, gemini	BNB+XRP	+\$489.1	+\$689.2 (77) +	Survived
cw2-30ds1	gemini	BNB+XRP	+\$383.8	+\$40.9 (81) +	Survived
cw2-7d1	fable, opus, deepseek, gemini	SOL+XRP	+\$299.1	+\$817.7 (258) +	Survived
cw2-7d2	gpt, deepseek, gemini, qwen	BNB+XRP	+\$223.7	+\$390.0 (77) +	Survived
cw2-pt1	fable, deepseek, grok [remap]	SOL	+\$698.3	+\$1,470.1 (147)	Survived
cw2-pt2	opus	XRP	+\$543.7	-\$633.1 (30) +	Faded
cw1-30d1	fable, deepseek, opus, qwen	ETH+SOL	+\$868.4	+\$581.2 (129)	Survived
cw1-30d2	gemini, gpt, deepseek	SOL+XRP+ETH	+\$590.6	+\$668.1 (99)	Survived
cw1-30d3	fable, deepseek, opus, qwen	BTC+ETH+SOL	+\$580.5	+\$160.0 (185)	Survived
cw1-30ds1	opus	SOL+XRP	+\$442.8	+\$999.1 (259)	Survived
cw1-30ds2	grok [remap]	SOL+XRP	+\$262.9	+\$215.1 (187)	Survived
cw1-30ds3	deepseek	XRP+BNB	+\$119.0	-\$171.3 (222) +	Faded
cw1-7d1	fable, deepseek, opus, gpt	SOL+XRP	+\$439.2	+\$1,202.3 (336) +	Survived
cw1-7d2	fable, qwen, gemini, opus	XRP+BNB	+\$429.5	-\$581.9 (17) +	Faded
cw1-7d3	gpt, opus	BNB+XRP+SOL	+\$398.6	-\$702.7 (306) +	Faded
cw1-7ds1	gemini	XRP+BNB	+\$417.8	-\$626.9 (145) +	Faded
cw1-7ds2	gpt	XRP+BNB	+\$366.3	-\$299.6 (260) +	Faded
cw1-7ds3	qwen	XRP+BNB	+\$267.0	-\$605.7 (19) +	Faded

n in parentheses; * = fewer than 10 closed trades (insufficient cell); + = non-zero skipped_nofunds (entries the engine could not fund -- the printed economics are then not the economics that ran). Home = the number the card's origin issue printed; the live presets have no home issue. [conc] = stake above this issue's cap for its universe size; [remap] = grok-4.5 replayed on the grok-4.6 lineage. Verdict rule: Survived = net-positive on the calendar quarter. Survivors are tinted green.

Counted: 30 of 56 past cards are net-positive on the calendar quarter (calendar August: 20 of 26 in Monthly Config Watch #1; calendar September: 17 of 53 in Monthly Config Watch #2) -- by cohort 1 of 3 Monthly-2 cards, 1 of 5 issue-8 cards, 3 of 5 issue-7 cards, 2 of 5 issue-6 cards, 1 of 4 issue-5 cards, 3 of 5 issue-4 cards, 3 of 3 Monthly-1 cards, 3 of 6 issue-3 cards, 7 of 8 issue-2 cards, 6 of 12 issue-1 cards. Config Watch #9 read the same 56 cards on two windows that run past this one's end (3 of its fresh week's 7 days and 24 of its four weeks' 28 are quarter days): 25 of 56 on its fresh week, 27 of 56 on its four weeks, 17 surviving both. 9 of the 17 of those survivors are net-positive here, and so are 21 of its 39 faded cards. The best replay is cw7-wk1 at +\$2,271.8 on 427 closed trades -- the configuration of this issue's card #1 -- and the worst cw1-7d3 at -\$702.7 on 306. 49 of the 56 replay rows carry unfunded entries (up to 522 on cw5-wk1), against 45 of 53 on Monthly Config Watch #2's calendar September (up to 135) and 49 of 56 on Config Watch #9's four weeks (up to 157). Live presets: 4 of 4 net-positive on the quarter (calendar August: 4/4; calendar September: 0/4), 0 of 4 survived both of Config Watch #9's windows.

The four live presets lead the table by standing rule. On the quarter: sol +\$2,766.4 on 110 closed, xrp +\$1,666.6 on 27 closed, bc2 +\$1,460.6 on 58 closed, 4h4 +\$987.9 on 117 closed -- all four net-positive, where Config Watch #9 recorded 0/4 on its fresh week, Monthly Config Watch #1 4/4 on calendar August and Monthly Config Watch #2 0/4 on calendar September. sol and xrp are marked [conc]: both run a \$900 stake on a single coin, twice the \$450 cap the stake table sets for a 1-coin universe, so their printed economics assume funding this issue's own rule would not grant. sol at +\$2,766.4 is above the search grid's quarter ceiling (+\$2,271.8): replay rows are not in the search grid. The bc2 row is also this issue's hand-tuned anchor (section below) and its +\$1,460.6 is the same simulation on both sides of the report.

Origin re-runs (drift baseline). Re-simulating the Monthly-2 cards on their OWN origin window in this collect lands within -\$436.0..\$0.0 of the numbers Monthly Config Watch #2 printed: 2 of the 3 reproduce to the cent. The issue-8 cards drift -\$162.0..\$0.0 (3 of 5 to the cent). The largest gap is cw5-m301 (+\$1,717.1 printed, +\$1,219.8 on the re-run). Snapshot principle: no past issue is restated; the full drift table (nine cohorts; the issue-1 cards carry no origin re-run) is in the open data.

What this column is. 30 of 56 past cards are net-positive over 92 days (78 with data). Three measured facts bound the reading. The window contains the cards' own: every one of the ten cohorts searched on days inside it

(56 rows). The rows Config Watch #9 dropped on its two-window rule did not all lose here: 21 of its 39 faded cards are net-positive on the quarter, and 9 of its 17 survivors. And the fresh search moved with them: the search grid's quarter ceiling is +\$2,271.8 against +\$1,618.4 on calendar August, +\$1,041.5 on calendar September and +\$250.4 on Config Watch #9's fresh week, all from the same engine and the same criteria. The quarter is a longer window, not a cleaner one.

How much of the quarter is in-sample, counted. Every cohort's search windows lie inside the quarter, so no row here is a forward test in the weekly sense. Of the quarter's 78 data days, the share outside a cohort's own search windows runs from 60.3% (Monthly-1) to 64.1% (issue-8, issue-7, issue-6, issue-5, issue-4, issue-3 and issue-1); the data days after a cohort's publication day -- the forward part of the quarter for that cohort -- run from 48 (issue-1, published Aug 13) to 0 (Monthly-2 and issue-8, published Oct 3 and Oct 1, after the quarter ended). By cohort, share outside / days after publication: Monthly-2 61.5% / 0; issue-8 64.1% / 0; issue-7 64.1% / 7; issue-6 64.1% / 14; issue-5 64.1% / 21; issue-4 64.1% / 26; Monthly-1 60.3% / 26; issue-3 64.1% / 35; issue-2 61.5% / 42; issue-1 64.1% / 48. The table is in the open data.

Survival curve by generation

Each published generation of cards, read on the fresh weeks that followed it. A cell counts that generation's cards that finished net-positive on the k-th fresh week after the issue's own window edge, out of all of its cards, taken from the issue that ran that replay; the last column is this issue's calendar-quarter replay of the same cards. Nothing is re-simulated for this table: every number in it is already in a published open-data file, and this issue only lines them up. The chart under TL;DR draws the ten card rows of this table that have a fresh week, as shares; it is printed once, there.

Generation	Cards	+1 week	+2 weeks	+3 weeks	+4 weeks	+5 weeks	+6 weeks	+7 weeks	+8 weeks	Quarter
Config Watch #1	12	6/12	11/12	1/12	2/12	3/12	11/12	1/12	1/12	6/12
Config Watch #2	8	8/8	4/8	1/8	2/8	5/8	6/8	2/8	--	7/8
Config Watch #3	6	2/6	2/6	2/6	5/6	2/6	2/6	--	--	3/6
Config Watch #4	5	1/5	0/5	4/5	4/5	3/5	--	--	--	3/5
Monthly Config Watch #1	3	2/3	0/3	1/3	3/3	2/3	--	--	--	3/3
Config Watch #5	4	1/4	3/4	2/4	3/4	--	--	--	--	1/4
Config Watch #6	5	4/5	2/5	3/5	--	--	--	--	--	2/5
Config Watch #7	5	4/5	4/5	--	--	--	--	--	--	3/5
Config Watch #8	5	2/5	--	--	--	--	--	--	--	1/5
Monthly Config Watch #2	3	3/3	--	--	--	--	--	--	--	1/3
Live presets	4	2/2	2/4	1/4	3/4	3/4	2/4	0/4	--	4/4

Cell by cell: every +k cell is the walk-forward of the issue whose fresh week it is -- Config Watch #2 (Aug 10-17, 2026), Config Watch #3 (Aug 17-24, 2026), Config Watch #4 (Aug 24-31, 2026), Config Watch #5 (Aug 31-Sep 7, 2026), Config Watch #6 (Sep 7-14, 2026), Config Watch #7 (Sep 14-21, 2026), Config Watch #8 (Sep 21-28, 2026), Config Watch #9 (Sep 28-Oct 5, 2026) -- and a generation's +1 week is the first of those weeks after its own window edge: issue-1 at #2, issue-2 at #3, issue-3 at #4, issue-4 and Monthly-1 at #5, issue-5 at #6, issue-6 at #7, issue-7 at #8, issue-8 and Monthly-2 at #9. The live presets are not a generation and have no publication date: their seven cells are the seven issues that replayed them -- Config Watch #3 replayed 2 of the 4, #4 to #9 all 4 -- and that row is not drawn on the chart. Quarter is this issue's window (Jul 1-Sep 30, 2026, one 92-day column -- there is no calendar-month column here) for every row; Monthly Config Watch #2 read the issue-1 row at 2/12, the issue-2 row at 0/8, the

issue-3 row at 0/6, the issue-4 row at 2/5, the Monthly-1 row at 1/3, the issue-5 row at 2/4, the issue-6 row at 4/5, the issue-7 row at 4/5, the issue-8 row at 2/5 and the presets row at 0/4 on calendar September. Config Watch #9's own cards are not a generation of this table: this collect does not replay them (walk-forward above). Two properties of the table to read with it: a column is not one calendar week across rows (issue-1's +1 week is Aug 10-17, 2026, Monthly-2's is Sep 28-Oct 5, 2026), and issue-1's +1 week overlaps its own search window (Aug 5-12, 2026) by 2 days, Monthly-1's overlaps calendar August by 1 day (Aug 31) and Monthly-2's overlaps calendar September by 3 days (Sep 28-30).

Survivors vs faded: structure

Config Watch #9 replayed 56 past cards on two fresh windows (Sep 28-Oct 5, 2026 and Sep 7-Oct 5, 2026) and its standing rule kept 17 of them (net-positive on both) and dropped 39. This table puts the two groups side by side on the traits recorded for every card: the frozen query (horizon/timeframe, council size, universe size, calibration, minimum confidence, table stake, membership) and the replay statistics on this issue's calendar quarter. Neither window decided alone: 25 cards cleared the week, 27 the four weeks and 17 both. 17 rows against 39 rows: this is a description of two groups, not a test, and no trait here is claimed to cause survival.

Trait	Survivors (17)	Faded (39)
FH/TF	4h/1h (11), 1d/1d (3), 1d/1d,4h (1), 4h/4h (1), 4h/4h,1h (1)	4h/1h (15), 1d/1d (11), 4h/4h (5), 4h/4h,1h (3), 1d/1d,4h (2), 1d/4h (2), 1h/1h (1)
Council size	4 models (9), 3 models (3), solo (3), 2 models (1), 5 models (1)	solo (14), 3 models (10), 4 models (10), 5 models (3), 2 models (2)
Universe size	2 coins (15), 5 coins (2)	2 coins (31), 3 coins (3), 1 coin (2), 4 coins (2), 5 coins (1)
Calibrate ON	5 of 17 (29.4%)	5 of 39 (12.8%)
Min confidence	60 (17)	60 (32), 0 (3), 50 (2), 55 (2)
Table stake	\$450 (14), \$200 (2), \$375 (1)	\$450 (35), \$300 (3), \$250 (1)
Median maxDD	80.2%	87.0%
Median win rate	63.9%	61.8%
Median smoothness R2	0.503	0.331
Model membership	gemini (15), opus (11), fable (9), qwen (7), deepseek (6), gpt (6), grok (1)	gemini (24), deepseek (18), opus (17), gpt (14), fable (12), qwen (11), grok (7)

Counts are cards, not simulations, and each card contributes once per row; the membership row counts a model once for every card it appears in, so it sums above the group size. Medians are over each group's replay on this issue's calendar quarter (Jul 1-Sep 30, 2026), and the R2 median is taken over the rows that have one (17 of 17 survivors, 39 of 39 faded: a replay with too few days carries no fit). On Config Watch #9's own windows the same medians read maxDD 10.5% against 24.1% and win rate 72.7% against 57.9% on its fresh week (Sep 28-Oct 5, 2026), and maxDD 32.6% against 63.3% and win rate 71.4% against 58.6% on its four weeks (Sep 7-Oct 5, 2026). On this issue's calendar quarter the same two groups read 9 of 17 and 21 of 39 net-positive.

Calibrate axis: does TP/SL calibration help a config?

Matched pairs -- identical knobs, calibration OFF vs ON -- measured on net PnL. This issue reports **two different measurements** and does not pool them: the whole stage-2 grid of the calendar quarter (every matched pair the search produced) and the board-leader slice (pairs whose OFF or ON leg reached the top of a board). They answer different questions.

Measurement	Pairs	Calibrate wins	Mean delta	Best delta	Worst delta
Whole stage-2 grid (this window)	3,497	56.1%	+\$88.3	+\$1,847.0	-\$2,100.0
-- mature rows only	2,181	58.2%	+\$51.4	+\$1,316.6	-\$2,100.0
Board-leader slice (top of board)	13	0.0%	-\$1,587.3	-\$644.6	-\$2,100.0

On this window the axis is positive in the population and negative at the top. Calibration improves 56.1% of the 3,497 matched pairs on the whole grid (mean +\$88.3, median +\$26.1) and 58.2% of the 2,181 mature pairs (mean +\$51.4). On the board-leader slice it improves none of the 13 pairs (mean -\$1,587.3, median -\$1,715.2): the OFF legs there are at a median of +\$1,878.1 and the ON legs at +\$356.8, and 9 of the 13 ON legs are net-positive. The same axis improved 37.9% of its pairs on calendar August (Monthly Config Watch #1), 65.6% on calendar September (Monthly Config Watch #2) and 68.6% on Config Watch #9's fresh week. Where the axis reached the

cards is visible above: 1 of the 2 council cards is a calibrated cell and the solo card is not, and 25 of the 45 finalists run it.

The open-data pack also carries a 400-pair export for this window. Its delta distribution is not the population's (median +\$736.6 against the grid's +\$26.1), so no claim in this section is computed from it; it is published for inspection only.

Secondary analysis -- solo-model top

Sidebar to the council narrative. Solo configs run inside the same pipeline and face the same gates: ≥ 2 coins, the maturity gate for the window (≥ 30 closed), one config per model.

Solo top -- QUARTER (Jul 1-Sep 30, 2026)

#	Model	Coins	FH/TF	Net	Net %	maxDD	Trades	WR
#1	gemini	SOL+XRP	4h/1h	+\$1,491.0	+149.1%	80.8%	443 *	60.9%
#	Ladder	BE	Shifts	Conf filter	Source	Entry	Positional	Stake
#1	50:40,100:60	on	TP -10 / SL +75	≥ 60 (each)	median	0/0	reenter	\$450

One solo card, not two. Every eligible row on the mature solo board is the same model (gemini), so the one-config-per-model rule admits exactly one. No single-coin row outranks it on that board this window. The card carries 46 unfunded entries (flagged * on Trades); the best mature solo row with no unfunded entry is a different config (opus on SOL, +\$949.9 on 171 closed).

Council vs solo. On this window the best solo sits **below** the best council (+\$1,491.0 against +\$2,271.8 on the mature boards), on 443 closed trades against 427, with an 80.8% drawdown against 74.3% and 46 unfunded entries against 23. Config Watch #9 measured the same comparison on its two windows: the solo above the council on the week, the council above the solo on the four weeks. Same in-sample caveats apply to both sides.

OPEN IN SANDBOX marketmania.ai/s/cwq1-s1

Each link opens this exact frozen window; results visible without sign-in (embedded share signature). The codes are minted by the publish script before this file goes live; every card's full frozen query signature is in the open-data JSON next to its simulation index.

Quality boards

Net PnL stays the primary board (continuity with the weekly series); four selection views sit on top of it at zero additional simulations. **Calmar-like** = $\text{net} / \max(\text{maxDD}, 1.0)$, positive net only; **WR board** = win rate among mature configs (≥ 30 closed); **Smoothness** = R2 of a linear fit through the daily closed-equity series, mature rows with net > 0 and at least 5 days, ties broken by ulcer index; **Quality composite** = mean rank over net, Calmar-like, win rate and R2, lower is better. R2 and the ulcer index are computed by the collect on the engine's daily settled-equity series: one point per day that settled a trade, from the first such day (the \$1,000 start is not a point), index-spaced, the ulcer index over the running peak of those points -- so days without a settlement are compressed and a loss on the first settled day does not enter the ulcer index. Top-3 of each board below; all 15 rows of all four boards are in the open-data JSON.

Board	Config	Coins	FH/TF	Net	WR	maxDD	Cls	R2
QUARTER Calmar	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$944.2	95.7%	12.4%	47	0.985
QUARTER Calmar	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$629.4	95.7%	8.3%	47	0.985
QUARTER Calmar	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$786.8	95.7%	10.4%	47	0.985
QUARTER WR	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$944.2	95.7%	12.4%	47	0.985
QUARTER WR	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$786.8	95.7%	10.4%	47	0.985
QUARTER WR	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$629.4	95.7%	8.3%	47	0.985
QUARTER Smoothness	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$629.4	95.7%	8.3%	47	0.985
QUARTER Smoothness	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$786.8	95.7%	10.4%	47	0.985
QUARTER Smoothness	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$944.2	95.7%	12.4%	47	0.985
QUARTER Quality	fable, opus, deepseek, gemini [cal]	SOL+XRP	1d/1d	+\$1,361.5	90.4%	18.4%	73	0.965
QUARTER Quality	fable, opus, deepseek, gemini [cal]	SOL+XRP	1d/1d	+\$1,344.8	90.3%	18.4%	72	0.966
QUARTER Quality	fable, deepseek, gemini [cal]	SOL+XRP	1d/1d,4h	+\$944.2	95.7%	12.4%	47	0.985

[cal] = calibrate=ON cell; the board column names the window. The Calmar board is not degenerate this window -- its top row sits at 12.4% maxDD, above the 1.0 floor of $\text{net} / \max(\text{maxDD}, 1.0)$, so it ranks by the ratio. None of the 12 printed rows is a single-coin config, and all 12 are calibrated. The pool is 7,126 mature configs at the ≥ 30 gate, against 3,265 on calendar August and 2,967 on calendar September at the monthly ≥ 20 gate.

This issue runs one window, so the cross-window stability board is not computed; the weekly and monthly issues keep it (Config Watch #9: 7 configurations on WK and M30; Monthly Config Watch #2: 22 on M30 and MONTH).

Hand-tuned baseline vs the grid

Standing section, carried from the weekly series. One hand-tuned configuration is treated as the reference the search has to beat -- not on net PnL, where any large search wins by construction, but on the three properties a configuration is actually tuned for: **win rate, shallow drawdown and a straight equity line**. The reference (the **hand-tuned anchor**) is the live preset 1d_3_llm_BC2_v02, a 3-model 1d council on SOL+XRP at the \$450 two-coin stake; it is distinct from the pre-search hand-tuned set A/B/C carried in Baselines below. Its replay row is in the walk-forward table above; here it is the yardstick.

Beat rule (frozen with the criteria): a config beats the anchor only if it does so on all three at once -- win rate higher, maxDD lower and smoothness R2 higher -- among mature, non-concentrated search rows. Net PnL is not part of the rule; it is reported next to the winners so the cost of the improvement is visible.

The anchor on this window

Window	Config	Coins	FH/TF	Net	WR	maxDD	R2	Ulcer	Cls
QUARTER	Hand-tuned anchor (1d_3_llm_BC2_v02)	SOL+XRP	1d/1d	+\$1,460.6	77.6%	62.9%	0.562	11.26	58

343 of the 7,126 mature configs (4.8%) clear all three bars at once. The anchor's own quarter is +\$1,460.6 net at 77.6% win rate with a 62.9% drawdown, an R2 of 0.562 and 58 closed trades. Of the ten best of those 343 by net, 1 beats the anchor on net PnL as well; the other 9 buy the three quality bars with a smaller number, and the anchor's +\$1,460.6 is above every one of them. 5 of the ten are calibrated cells, every one is a council, their stakes are \$300, \$375 and \$450 and their universes are 2 coins. On calendar August (Monthly Config Watch #1) the same rule admitted 12 of 3,265 configs (0.4%), and on calendar September (Monthly Config Watch #2) 679 of 2,967 (22.9%), so the share moves with the window: read it as one field at a time, not as a property of the anchor.

Top-10 by net among the 343 configs that beat the anchor on win rate, maxDD and smoothness at once

#	Config	Coins	FH/TF	Stake	Net	WR	maxDD	R2
#1	gemini, gpt, deepseek	SOL+XRP	1d/1d	\$375	+\$1,685.4	78.0%	60.9%	0.847
#2	gemini, gpt, deepseek	SOL+XRP	1d/1d	\$450	+\$1,455.1	79.2%	55.2%	0.764
#3	gemini, gpt, deepseek	SOL+XRP	1d/1d	\$450	+\$1,389.4	78.0%	61.9%	0.571
#4	fable, opus, deepseek, gemini [cal]	SOL+XRP	1d/1d	\$450	+\$1,361.5	90.4%	18.4%	0.965
#5	gemini, gpt, deepseek [cal]	SOL+XRP	1d/1d	\$450	+\$1,355.0	94.0%	27.6%	0.941
#6	gemini, gpt, deepseek	SOL+XRP	1d/1d	\$300	+\$1,348.3	78.0%	48.7%	0.847
#7	fable, opus, deepseek, gemini [cal]	SOL+XRP	1d/1d	\$450	+\$1,344.8	90.3%	18.4%	0.966
#8	gemini, gpt, deepseek [cal]	SOL+XRP	1d/1d	\$450	+\$1,323.0	89.7%	43.5%	0.796
#9	opus, gpt, deepseek, gemini, qwen [cal]	SOL+XRP	1d/1d	\$450	+\$1,269.1	89.8%	51.7%	0.881
#10	fable, gpt, gemini, qwen, grok	SOL+XRP	4h/4h,1h	\$450	+\$1,212.9	78.2%	49.2%	0.813

Anchor passport (verbatim, as stored): members=gemini-3.1-pro,gpt-5.6-sol,deepseek-v4-pro tokens=SOL,XRP fh=1d tf=1d conf_mode=each min_conf=60 max_sideways=2 max_diff_side=0 tp_sl_source=nearest trade_mode=position same_side=update opposite=reverse no_signal=hold time_stop=3x steps=40:80,100:20 steps_on=1 be=1 tp_shift=30 sl_shift=50 deposit=1000 lev=10 refill=0 stake=450 calibrate=- variant=v0

Per-ticker bests

Which coin was extractable over the quarter, and by what. Single-coin results are excluded from the council cards by rule (idiosyncratic-coin risk) and reported here, never ranked against the cards. The dedicated per-coin branch contributed 585 sims -- a compacted grid, one best per coin per window; the branch tag in brackets says which kind won.

Coin	Best config (branch)	FH/TF	Net (trades)
BTC	deepseek [solo]	1d/1d	+\$154.6 (2) *
ETH	opus, deepseek, qwen [council]	1d/1d	+\$498.5 (39)
SOL	fable, deepseek, grok [council]	4h/1h	+\$1,470.1 (147)
BNB	opus, gpt, deepseek, gemini, qwen [council]	1d/1d	+\$809.2 (31)
XRP	fable, deepseek, gemini [council]	1d/1d	+\$1,147.0 (32)

n in parentheses; * = fewer than 10 closed trades (thin -- direction only). 1 of the five winners runs calibrate ON and 4 of the five are councils.

The reads. Every coin is net-positive over the quarter and 1 row is thin (under 10 closed trades); the five run 2 to 147 closed trades, SOL leading at +\$1,470.1 and BTC lowest at +\$154.6. 1 of the five sits below the ≥ 30 -trade maturity gate the council boards apply (BTC 2), so that row is direction with a small sample rather than a measured edge. The 1d horizon wins 4 of the five coins. **Winner's curse applies to this whole table:** each row is the maximum of a per-coin grid, biased high by selection alone, and none has an out-of-sample read until the next issue.

Baselines

Three independent reference points, none of them search output, re-simulated fresh on this issue's frozen window inside the same collect (engine 1.1): the engine solo default for each of the 7 models, the default council-of-7, and the pre-search hand-tuned set A/B/C carried since Config Watch #1. The table shows the universe-aware **table-stake** run -- this issue's canonical mode (\$200 on 5 coins for the defaults; A and B already at their \$300 3-coin cap; C moves \$400 -> \$450 on 2 coins) -- with the best and worst of the seven solo defaults; all eleven rows and both stake modes are in the open data.

Baseline	Net PnL	Trades	WR	maxDD
Best solo default (gemini-3.1-pro)	-\$229.5 *	1079	53%	24.2%
Worst solo default (deepseek-v4-pro)	-\$397.0 *	2312	56%	40.5%
Council-of-7 default	-\$161.7 *	1165	60%	17.3%
Hand-tuned config A (pre-search)	-\$326.8 *	61	67%	148.5%
Hand-tuned config B (pre-search)	+\$112.4 *	392	65%	84.8%
Hand-tuned config C (pre-search)	+\$581.2	129	66%	68.5%

* = non-zero skipped_nofunds at table stake. On the 5-coin defaults those counts are the stake table doing its job: at \$200 a \$1,000 deposit funds at most five concurrent positions (1,461 to 3,177 skipped each over the quarter). The other five solo defaults run between the two printed rows, and both stake modes carry the same sign for every one of the eleven rows -- **the conclusions do not change between modes**. Hand-tuned C moves \$400 -> \$450 on 2 coins and reads +\$516.6 at its own stake against +\$581.2 at the table stake.

Over the quarter the standing ordering holds in one half only. The seven solo defaults run -\$229.5 to -\$397.0 at table stake and the council-of-7 -\$161.7 -- every default lost money at table stake; the three pre-search hand-tuned configs read -\$326.8 / +\$112.4 / +\$581.2, and the fresh search's +\$2,271.8 is above the best reference point in the table (+\$581.2). Search > hand-tuned > defaults is only partly what this quarter measures: the search maximum is above every reference point; a default is at or above a hand-tuned config. The live presets are a separate row of evidence: all four are net-positive on the quarter (walk-forward table above), and two of them run stakes above this issue's cap.

Live presets & preset neighbourhoods

Four hand-tuned configurations trade live on the platform (not search output, distinct from the pre-search A/B/C set): preset bc2 = '1d_3_llm_BC2_v02'; preset 4h4 = '4h_4_llm_v02'; preset sol = 'Sol'; preset xrp = 'XRP'. Their replay rows lead the walk-forward table; passports and window detail are in the open data. The neighbourhood grid asks a narrower question: does a one-knob neighbour beat the preset as configured? Two of the four have a local grid this issue (53 cells on this window); sol and xrp do not (0 cells -- the local grid is not built for them, so no neighbour claim is made about either).

Live preset	As configured	Best neighbour in local grid	Cells	Verdict
preset bc2	+\$1,460.6	+\$2,022.5 (1d/1d, SOL+XRP, 50 cls)	34	neighbour +\$561.8
preset 4h4	+\$987.9	+\$1,027.4 (4h/1h, ETH+SOL, 117 cls)	19	neighbour +\$39.5
preset sol	no local grid this issue	--	0	--
preset xrp	no local grid this issue	--	0	--

In **both** preset cells with a grid, a one-knob neighbour beat the configuration as it is actually running on this window: bc2 by +\$561.8 (+\$1,460.6 -> +\$2,022.5; tp_sl_source nearest -> average); 4h4 by +\$39.5 (+\$987.9 -> +\$1,027.4; max_diff_side 0 -> 1). Cell counts are small (34 and 19) and every neighbour is an in-sample maximum of its own little grid.

Patterns: what the finalists look like

Knob modes across this window's unique finalists (mature net top-10 per branch plus all four quality boards, deduplicated by full signature; 45 unique configs, 35 of them councils and 10 solo).

Knob	QUARTER finalists mode
Forecast horizon	1d (32/45)
Ladder shape (steps)	50:40,100:60 (40/45)
Break-even stop (BE)	on (45/45)
Min confidence	60 (40/45)
SL shift	75 (40/45)
TP shift	-10 (40/45)
TP/SL source	median (32/45)
Calibrate ON	25/45
Universe	SOL+XRP (42/45)

The mechanical core: 50:40,100:60 ladder (40/45), break-even on (45/45), conf 60 (40/45), SL +75 / TP -10 (40/45), median source (32/45) -- Monthly Config Watch #1 and #2 printed 50:40,100:60 / on / 60 / 75 / -10 / farthest for calendar August and September. The horizon mode is **1d** (32/45; calendar August: 1d (46/62); calendar September: 4h (33/49)). **Calibrate=ON** runs in 25 of the 45 finalists, against 37 of 62 on calendar August and 31 of 49 on calendar September. The universe is SOL+XRP in 42 of the 45 finalists, and 2 of the 45 are single-coin rows. Membership across the 35 council finalists: gemini 35, deepseek 31, fable 30, opus 19, grok 10, gpt 7, qwen 2. Standing caveats: refinement seeds from the same grid winners, so part of the convergence is search-design echo, and a knob core that produced 17-of-56 survival on Config Watch #9's two windows and 30-of-56 net-positive on this quarter is describing windows, not a recommendation.

Liquidation note

Liquidation is modeled by the engine at the fixed 10x leverage used throughout. Across the published cards, stop-loss placement stays inside the distance that would approach the liquidation threshold; maxDD is printed on every config so realized risk is visible directly. On this window the published cards draw down 74.3% (council #1), 17.8% (council #2) and 80.8% (the solo card), and the replay table is deeper: 58 of the 60 replayed rows drew down more than 40% over the quarter and 56 past 55%. A high headline net and a survivable path are different claims, and so are a shallow in-sample drawdown and a shallow one next quarter.

Watch amendments (this issue)

Amendment 3 -- smoothness and a quality composite (standing). Every simulation stores its daily closed-equity series and the three numbers read off it (n_days , $smooth_r2$, $ulcer$); the smoothness board ranks mature rows with $net > 0$ and at least 5 days by $R2$ with $ulcer$ as tie-break, and the quality composite ranks by the mean of the net, Calmar-like, win-rate and $R2$ ranks. Both still cost zero additional simulations and the headline board is still net PnL.

Amendment 4 -- the monthly split (standing). The weekly Config Watch keeps two windows, the calendar week (WK) and the four calendar weeks that end with it (M30); the calendar-month board belongs to the Monthly Config Watch line. This issue prints no weekly or monthly board: its one window is the calendar quarter (Jul 1-Sep 30, 2026), and every weekly or monthly number in it is read from the issue that published it.

Amendment 5 -- the standing OOS board grows with every issue (standing). Every card this series publishes joins the walk-forward table and stays on it: this issue replays the 56 cards Config Watch #9 replayed -- Monthly Config Watch #2's 3, issue #8's 5, issue #7's 5, issue #6's 5, issue #5's 4, issue #4's 5, Monthly Config Watch #1's 3, issue #3's 6, issue #2's 8 and issue #1's 12 -- plus the four live presets, on the calendar quarter, under the standing rule with the one window this issue has (see Walk-forward above). Config Watch #9's own five linked cards are not in this collect.

All three amendments are standing, provisional and scoped to this series. **Formalization is slated for methodology v1.2**; until then this issue runs on v1.1 plus the Config Watch amendments approved 12.08 (hash e66c7e8c864a2233).

Limitations

- **One calendar quarter, one regime, partial data.** Every number here is a single 92-day window (Jul 1-Sep 30, 2026) with 78 days of platform data (the first forecast slot is Jul 15 18:01) on a single regime; the month figures set it beside the two monthly issues inside it (calendar August, Monthly Config Watch #1; calendar September, Monthly Config Watch #2), and no past issue is restated. Trades settle past the window edge at their horizon; open = a position still open when the collect ran (Oct 9, 07:08 UTC, 199 hours after the window edge): 0 on the three published cards, none of the 60 replay cells (60 rows x 1 window) carries one -- listed in the open data.
- **This quarter contains the monthly windows, so it is not independent of them.** 61 of its 78 data days lie in the two published monthly issues' months it is read against (Monthly Config Watch #1 Aug 1-31, 2026; Monthly Config Watch #2 Sep 1-30, 2026); Jul 15-31 lie in neither, and it holds 3 of Config Watch #9's 7 fresh-week days and 24 of its 28 four-week days; the quarter, the months and the weeks are readings of overlapping data, and agreement between them is partly arithmetic.
- **The replay column is decay context, not out-of-sample.** The quarter contains the search windows of all ten cohorts, so all 56 replayed cards are partly scored on the days they were selected on (the share of the quarter's data days outside a cohort's own windows runs from 60.3% to 64.1%). Publication dates bound it too: 48 of the 56 cards were published inside the quarter (issue-7 on Sep 23, issue-6 on Sep 16, issue-5 on Sep 9, issue-4 and Monthly-1 on Sep 4, issue-3 on Aug 26, issue-2 on Aug 19 and issue-1 on Aug 13), so for them the quarter's days before publication are not a forward test of a published card, and 8 were published after it ended (Monthly-2 on Oct 3 and issue-8 on Oct 1), so none of its days is; no card was published before Jul 1. The 56 cards are 52 distinct configurations window aside ($cwm2-m1 = cw8-m302$; $cwm2-m2 = cw8-m301$; $cwm1-s1 = cw3-s2$; $cw3-m301 = cw3-wk1$). The four live presets are not search output and have no home issue, and 2 of them run stakes above this issue's cap for their universe size ([$conc$] rows) -- their printed economics assume funding the current rule would not grant. Config Watch #9's five linked cards are not replayed here (walk-forward above).
- **The survival curve mixes window lengths and calendar weeks.** Each +kw cell is a 7-day window from a different issue and a different week of August, September and October, the Quarter column is a 92-day window with 78 days of data, and the live-preset row has no publication date at all. Read down a column, not across a row, and not as a decay rate.
- **Stability is defined inside the run that computes it.** The board needs two windows of one run; this issue's collect has one, so no stability board is printed or compared here -- Config Watch #9 reads its board across WK and M30 and Monthly Config Watch #2 across M30 and MONTH, and neither is a reading of the quarter.
- **Research-to-date counter is pinned, and it is read at this issue's cutoff.** The counter below is pinned from wave 3 onward and prints 59,755 resolved forecasts at the quarterly cutoff (Thu Oct 1, 2026, 16:00 UTC), where Config Watch #9 printed 62,260 at its own cutoff (Oct 5, 2026 16:00 UTC) and Monthly Config Watch #2 59,755

at the same cutoff: the same counter at different cutoffs, not separate measurements. The model line counts 7 current models with earlier versions folded into their successors' lineage; inside this quarter the collect's lineage probe finds 10 raw model IDs, 3 of them earlier versions spliced into a current lineage (claude-opus-4.8 until Jul 30, qwen-3.7-max until Aug 5 and grok-4.5 until Aug 24), so a council that names one of those models runs on the earlier version before that date.

- **Multiple testing, winner's curse, unfunded entries and thin cells.** 10,429 simulations in this collect; at this scale some winners are expected from chance alone, and every card, board row and per-ticker best is the maximum of a search, biased high by selection. On this window 0 replay rows are under 10 closed trades and 1 of the five per-ticker winners is under the ≥ 30 -trade maturity gate; 49 of the 56 past-card rows carry unfunded entries, and so council card #1 (23) and the solo card (46) do: their printed economics are not the economics that ran. Antidotes: axes frozen before launch, the walk-forward table, independent baselines, in-sample labeling; no formal correction yet. Research output, not financial advice.

Counters & lineage

Quarter collect (one window: the calendar quarter, printed whole)

Stage	Sims	Stage	Sims
Scan (s1)	476	Dedicated per-ticker (s6t+s6tc)	585
Systematic grid (s2)	9,028	OOS replays (s4cwm2+s4cwm2o+s4cw8+s4cw8o+s4cw7+s4cw7o+s4cw6+s4cw6o+s4cw5+s4cw5o+s4cw4+s4cw4o+s4cwm1+s4cwm1o+s4cw3+s4cw3o+s4cw2+s4cw2o+s4cw1+s4pre)	104
Refinement: ladder + stake sweep (s2b+s3b)	123	Total	10,429
Preset neighbourhood grid (s2p)	53	Errors / zero-trade	0 / 951
Per-ticker probes (s5t)	60	Engine	1.1

Scale note. This issue's collect is a run of its own: 10,429 simulations over its one window, against 10,889 in Monthly Config Watch #1's collect and 11,075 in Monthly Config Watch #2's. Config Watch #9 ran two separate collects (5,626 weekly + 11,030 night = 16,656), which this issue does not use. The stage-2 grid is 9,028 cells here, 4,176 on Monthly Config Watch #2's calendar month and 4,240 on Config Watch #9's fresh week. The replay branch is 60 rows -- 56 past cards and 4 live presets -- for 104 simulations, because every row is replayed on the one window of this collect and the cards of nine cohorts are re-run on their origin windows. The selection boards and the survival, structure and monthly tables cost zero simulations.

No extras pass. Equity series, the grid histogram (0 extra sims), the baselines and the calibrate pack are produced by the same collect that wrote the boards. The baseline runs are verification, outside the search total: 22 sims in this collect (11 configs x 1 window x 2 stake modes), all of them on the calendar quarter. 0 equity replays were needed.

Lineage probe: forecasts per raw model ID inside the quarter (Jul 1-Sep 30, 2026)

Raw model ID	Forecasts	First slot	Last slot	Lineage
claude-fable-5	19,009	Jul 15 18:01	Sep 30 23:01	current
claude-opus-4.8	6,315	Jul 15 18:01	Jul 30 08:01	folded into claude-opus-5
claude-opus-5	12,695	Jul 30 08:01	Sep 30 23:01	current
deepseek-v4-pro	19,007	Jul 15 18:01	Sep 30 23:01	current
gemini-3.1-pro	19,006	Jul 15 18:01	Sep 30 23:01	current
gpt-5.6-sol	19,010	Jul 15 18:01	Sep 30 23:01	current
grok-4.5	11,128	Jul 15 18:01	Aug 24 09:01	folded into grok-4.6
grok-4.6	7,870	Aug 24 10:01	Sep 30 23:01	current
qwen-3.7-max	7,479	Jul 15 18:01	Aug 5 09:01	folded into qwen-3.8-max
qwen-3.8-max	11,520	Aug 5 10:01	Sep 30 23:01	current

One read-only count of the platform's stored forecasts per raw model ID with a slot inside the window, written by the collect itself. 10 raw IDs: the 7 current ones and 3 earlier versions, each folded into its successor's lineage (claude-opus-4.8 into claude-opus-5, qwen-3.7-max into qwen-3.8-max and grok-4.5 into grok-4.6). Lineage = the current ID the engine runs a council member on.

Source files: summary_quarter.json (the w309 collect) with grid_hist_quarter, equity_quarter, baselines_quarter, presets_recon_quarter and the ten oos_*_quarter packs beside it; the published open data of Config Watch #2 to #9, Monthly Config Watch #1 and #2 and the three quarterly reports of Oct 8 (with their metrics pack quarterly_metrics_2026-q3 for the research-to-date counter) for every number this issue quotes from them. Generated at: 2026-10-09T07:08:39Z · Methodology v1.1 + Config Watch amendments (approved 12.08), hash e66c7e8c864a2233.

Snapshot principle: this report is a frozen snapshot of the search taken at the generated-at timestamp above; numbers are not updated retroactively. Each issue re-runs the pipeline fresh over that issue's window.

Research to date

RESEARCH SNAPSHOT

THIS REPORT -- search effort: 10,429 simulations (0 errors) in one frozen collect, of which 104 walk-forward replays and origin re-runs · 22 verification sims (baseline runs inside the same collect; 0 equity replays needed)

MARKETMANIA RESEARCH TO DATE (as of cutoff Oct 1, 16:00 UTC): 59,755 resolved since Jul 11 · 7 models tracked (current line-up; earlier versions folded into their successors' lineage) · 5 assets · 5 horizons · hourly · 50 published reports

Counter pinned from wave 3 onward and read at this issue's cutoff (Thu Oct 1, 2026, 16:00 UTC), where Quarterly Consensus Watch #1, Quarterly Calibration #1, Quarterly Model Watch #1 and Monthly Config Watch #2 read the same 59,755 and Config Watch #9 read 62,260 at its own later cutoff; the model line counts the current line-up, with earlier versions folded into their successors' lineage. Quarterly Config Watch #1 is the 50th published report: Quarterly Consensus Watch #1, Quarterly Calibration #1 and Quarterly Model Watch #1 are 47, 48 and 49, and Config Watch #9 is 46.

Issue #1. Quarterly Config Watch is a living quarterly comparison: each issue appends one more calendar quarter of cards, walk-forward verdicts and survival data, and nothing in a past issue is restated. This issue opens the quarterly cadence on a partial quarter (Q3 2026, platform data from Jul 15) -- a best mature council of +\$2,271.8 and a best mature solo of +\$1,491.0, 30 of 56 published cards net-positive on the quarter, a survival curve that reaches +8 fresh weeks for the issue-1 generation, and a structural table of Config Watch #9's 17 survivors against its 39 faded cards; there is no prior quarter, and the two monthly issues inside it (calendar August and September) are the comparison. The calendar week, the four weeks and the calendar month stay in the weekly and monthly Config Watch.

What we're testing next

- **Quarterly Config Watch #2** -- the Q4 2026 calendar quarter (2026-10-01 to 2027-01-01) from the same engine, the same criteria and the same collect code, read against the monthly issues whose months lie in that quarter. It carries the first calendar-quarter replay rows for this issue's own cards.
- **This issue's cards on a fresh week.** The weekly Config Watch #10 is the first issue to replay the three cards of this issue, on its fresh week and beside every other card on the walk-forward board.
- **The survival curve extended.** Every weekly issue adds one fresh week to every generation already on the curve and puts a new generation on it; Quarterly #2 prints those columns for the Q4 2026 weeks, with the issue-9 generation on the curve for the first time.

Related research

Quarterly Consensus Watch #1 -- <https://marketmania.ai/research/reports/consensus-watch-quarterly-2026-q3.pdf>

Quarterly Calibration #1 -- <https://marketmania.ai/research/reports/calibration-quarterly-2026-q3.pdf>

Quarterly Model Watch #1 -- <https://marketmania.ai/research/reports/model-watch-quarterly-2026-q3.pdf>

Config Watch #9 -- <https://marketmania.ai/research/reports/config-watch-2026-09-30.pdf>

Monthly Config Watch #2 -- <https://marketmania.ai/research/reports/config-watch-monthly-2026-09.pdf>

The three quarterly reports published together on Oct 8 as one issue; this issue joins them on the same cadence. Config Watch #9 is the newest weekly issue: it carries the calendar week Sep 28-Oct 5, 2026 and the four-week window Sep 7-Oct 5, 2026 from collects of its own, and no board of either is reprinted here. Monthly Config Watch #2 is the newest monthly issue and, with Monthly Config Watch #1, the source of every month-against-quarter number.

CITE THIS REPORT

```
@techreport{mm_configwatch_quarterly_2026q3,
  title      = {Quarterly Config Watch #1},
  author     = {{MarketMania Research}},
  institution = {MarketMania},
  year       = {2026},
  month      = oct,
  day        = {9},
  type       = {Quarterly Research Report},
  series     = {Quarterly Config Watch},
  number     = {1},
  note       = {Methodology v1.1 + Config Watch amendments, hash e66c7e8c864a2233},
  url        = {https://marketmania.ai/research/reports/config-watch-quarterly-2026-q3.pdf}
}
```