

BTC Operating System Deep Test Report

BTC Operating System v1.1.1 • Updated 2026-06-01

Official links: [TradingView page](#) • [Support](#)

This report is the deeper buyer-facing BTC proof layer for the current public release. It tests whether the live reserve-core workflow earns its place over blind weekly auto-buy under matched assumptions.

1. Current v1.1.1 scope audit

This report is the deeper buyer-facing BTC proof layer for the current public release. It is built to answer one practical question: does the live reserve-core workflow justify itself over blind weekly auto-buy under matched assumptions, for real buyers with real contribution sizes and a real weekly process?

The current BTC release remains a weekly reserve-core allocation system. The public proof cut keeps the same operating job, the same dashboard-first weekly workflow, and the same matched-control benchmark standard while improving version discipline and customer-legibility.

2. What the benchmark actually measures

BTC is not sold on generic “discipline.” It is sold on whether the engine improves Bitcoin-equivalent outcomes relative to the weekly auto-buy control under the same buyer inputs.

Published scoreboard row	Meaning
AZRO BTC-eq	Ending BTC plus ending cash translated back into BTC-equivalent at the reference weekly close.
Auto BTC-eq	The matched simple weekly auto-buy control under the same buyer inputs.
× vs Auto	AZRO BTC-eq divided by Auto BTC-eq.
Extra sats	BTC-equivalent lead versus the control shown in satoshis.
Avg % ahead of auto-buy	Average annualized lead versus the matched control across the chapter.

Control basis: Same Start date, same Weekly budget, same Cash at Start date, same BTC already owned context, same fee handling, same weekly-close convention.

3. Current published BTC event-family evidence

The current public BTC evidence layer includes published event-family summaries, the matched-control on-chart scoreboard, a contribution-fit guide, and a chapter-continuity workflow. That is still not the future public weekly benchmark archive, but it is enough to judge whether the current release is commercially serious and benchmark-accountable.

- Audit dimensions covered in this customer cut:
- Matched-control benchmark logic: same chapter inputs on both sides so the comparison stays honest.
- Published event-family behavior: top-side major events, re-accumulation windows, and rare deep zones with forward medians.
- Contribution-fit economics: where the reserve-core engine is easier or harder to justify relative to DCA.
- Workflow realism: one calm weekly decision, one tracker-backed continuity layer, and one benchmark standard that stays visible after purchase.

Published BTC event family	Signals	Historical span	Median 13W	Median 26W	Median 52W
Top-side major events	15	2017-10-30 to 2025-10-13	-23.4%	-22.8%	-41.6%
Bottom-side re-accumulation	10	2015-09-14 to 2023-08-28	+16.3%	+66.2%	+124.4%
Rare deep zone	3	2015-09-21 to 2022-11-07	+33.6%	+83.1%	+127.3%

4. How to read BTC versus the auto-buy baseline

- AZRO can only earn its place three ways: better buying in weak windows, earlier defense in expensive or risky windows, and cleaner continuity over time.
- The auto-buy baseline remains the honest baseline and should stay public.
- The auto-buy baseline may still be enough in clean grind-up regimes where almost every week should simply be bought.
- The auto-buy baseline may still be enough when contribution size is too small for optimization to matter much after frictions.
- Visible lag is part of honest proof. If the engine lags, the benchmark should show it clearly.

5. Illustrative benchmark situations

- DCA can win in a clean grind-up where almost every week should simply be bought and there is little reason to defend early.
- AZRO can earn its place when weak windows allow materially better buying, when overheated windows justify earlier defense, or when the buyer values lower weekly decision stress and chapter continuity enough to care about the control.
- The reserve-core question is not “can it ever win?” The reserve-core question is whether the engine earns its complexity often enough, and visibly enough, for your contribution size and discipline.

6. What a serious buyer should verify next

- Read the BTC Handbook for the hard standard, workflow rules, chapter continuity, and alert paths.
- Read the BTC Public Evidence Brief and BTC Proof Brief for the compressed public case.
- Use the Companion Tracker if you intend to run a real chapter and keep continuity auditable.
- Use the fit guide to judge whether your contribution size and workflow preference justify a managed reserve-core engine over simple recurring buying.

7. Current release conclusion

The current BTC public layer is strong enough to judge as a reserve-core system: it publishes the matched-control scoreboard on chart, publishes event-family evidence, keeps the workflow calm and weekly, and makes the benchmark standard explicit.

What this report proves now: the core offer is legible, benchmark-accountable, and commercially serious.

What it does not claim: that every window will outperform DCA or that the public customer pack is the full end-state trust layer. The public weekly proof layer still belongs on the site next.