The AfterFiat Operator & Investor Checklist

Jason St George · extracted from *AfterFiat v2.6* §25 · 2 September 2026 · afterfiat.xyz

> Everything before this checklist argues the thesis at full strength; a checklist is where you stop believing the argument and start checking it. Each question has a yes that strengthens the claim and a no that weakens or falsifies it.

1. Triad supply: Privacy, Proofs, Compute

Privacy

- Live, non-custodial privacy corridors with measurable anonymity sets?
- Shielded pools growing in depth and churn?
- Lawful-privacy patterns (viewing keys + receipts) actually used?

Proofs

- Robust proof supply for canonical workloads?
- VERIFYPRICE SLOs met consistently?
- Proof production diversified across hardware profiles and geographies?

Compute

- Meaningful amounts of useful compute being proven?
- Liquid market for Work Credits?
- Triad workloads tied to real-world demand?

2. Value capture & anti-bypass

The most important questions for whether the SoV thesis holds. If the answer to any is "yes," the system may be useful infrastructure but the SoV claim is weakened or falsified.

- Can users buy equivalent Privacy, Proofs, or Compute service **without the native asset**?
- Are fees **actually paid in the asset**, or do operators accept and immediately off-ramp alternatives?
- Are burns and collateral lockups **material** relative to total supply?
- Is there a visible **bypass channel** (hyperscaler-hosted provers, stablecoin-denominated services) capturing most demand without flowing through the asset?
- Does governance have the power to **redirect value away from holders** (eliminating burns, inflating supply, changing fee routing)?
- Is the **Wrapper Dominance Ratio** rising while native fee share, private settlement, and collateral lockups stagnate?
- Is **price** being cited as evidence of monetary adoption in place of native fee, burn, collateral, and settlement series?

3. Institutional holdability — with counterweights

- Custody support from regulated custodians? Legal classification clear enough for mandates?
- Asset auditable for institutional reporting? Liquidity deep enough for institutional-scale positions?
- A valuation policy allocators can apply (fee-flow models, comparable frameworks)?

Holdability is a virtue with a bill attached. Pair each answer with its market-realization counterweight:

- Does wrapper growth **seed** native usage, or **substitute** for it? Is the Wrapper–Native Growth Gap persistently positive?
- What share of supply sits under custodial control, and how concentrated is that custody?
- Is delta-adjusted synthetic exposure large relative to liquid free float?
- Do daily-reset leveraged products exist, and what is their net mechanical gain and Wrapper Recycling Ratio?
- Is exposure concentrated in a small number of dealers, and has the hedge mix shifted from physical to swaps to options?
- Is demand governed by a small number of common rebalancing policies?

4. Stack health: Layers 0–6

Layer 0.
Documented hardware profiles? Open hardware participation? Facility Capacity Receipts independently verifiable? Physical VERIFYPRICE within bounds?

Layer 1.
Resilient connectivity with alternative transports?

Layer 2.
Clients distributed via multiple channels? Reproducible, signed binaries?

Layer 3.
Pseudonymous credentials integrated? Non-custodial usage share visible?

Layer 4.
Canonical workloads documented? VERIFYPRICE within SLOs? Proof production decentralized?

Layer 5.
Non-custodial corridors healthy and diversified? Refund safety enforced?

Layer 6.
Public dashboards? Documented governance? Recent incident reports?

Market realization plane.
Is VerifyFlow published? Wrapper inventory, creations/redemptions, flow elasticity, concentration, and growth gap visible? Stress harness run and reported?

5. Telemetry honesty

Telemetry itself can be gamed:

- **Open data.** Raw metrics public? Independent teams can reproduce Verify* metrics?
- **Client diversity.** Metrics gathered by multiple operators?
- **Incentive alignment.** Any incentive to under-report problems?
- **History.** Historical series, or only recent snapshots?

6. Red flags and failure patterns

Closed hardware monoculture.
One TEE, one vendor, no sampling, no open profiles.

Opaque bridges.
"Magic multisig" bridges with no proofs, unclear jurisdiction, no incident history.

Foundation fiat.
Parameter or policy changes via blog post, no on-chain trace, no incident report.

Zombie corridors.
Privacy rails that haven't moved meaningful volume in months but still appear in marketing.

Proof theater.
"ZK" branding with no public VERIFYPRICE metrics, no canonical workloads, no commodity verifiers.

Governance theater.
Token votes with single-digit participation deciding fundamental parameters.

Price as proof.
Marketing citing market cap, wrapper AUM, or trailing returns as evidence while native fee, burn, collateral, and settlement series are absent or flat — financialization narrated as adoption.

> A system with **no telemetry** is untrustworthy. A system with **one vendor's telemetry** is fragile. A system with **multi-source, reproducible telemetry** has a shot at being money.

Extracted from *AfterFiat: The Load-Bearing Thesis*, v2.6, §25 (Operator & Investor Checklist). The full record — including every metric referenced above with its measurement contract — lives at `afterfiat.xyz`. © 2026 Jason St George, CC BY 4.0.