

Next Generation Stores of Value

Privacy, Proofs, Compute

*A Conditional Monetary Thesis
for Verifiable Digital Infrastructure*

Jason St George

A cypherpunk monetary stack: privacy by default, proofs by construction, and compute that only gets paid when anyone can verify it.

Version 2.5

2026

Rights & License

© 2026 Jason St George.

This work (text and original figures) is licensed under **Creative Commons Attribution 4.0 International (CC BY 4.0)**. You may share and adapt the material for any purpose, including commercially, provided you give appropriate credit, provide a link to the license, and indicate if changes were made.

License: <https://creativecommons.org/licenses/by/4.0/>

Trademarks: “Next Generation Stores of Value,” “VerifyPrice,” “VerifyReach,” “VerifySettle,” “VerifyFlow,” “Work Credits,” and related marks are trademarks/service marks of the authors and are not licensed under CC BY 4.0. No endorsement implied.

Third-party materials: All third-party trademarks, images, and excerpts are the property of their respective owners and are used under fair use or by permission where applicable; they are not covered by the CC license unless explicitly stated.

How to cite:

Jason St George. *Next Generation Stores of Value: Privacy, Proofs, Compute*. Version 2.5. Licensed CC BY 4.0.

Abstract

Under sustained administrative and financial repression, a bearer asset may earn monetary premium if holding it preserves private settlement, portable proof, and access to verified compute after ordinary legal, custodial, and platform substitutes weaken. That premium is conditional: the full service path must remain stress-deliverable through independently substitutable physical routes; demand must not bypass the asset; a persistent, self-custodied, loss-bearing holder constituency must absorb residual financial risk; and the duration-bearing credit required to build the infrastructure must remain explicitly separate from the monetary object. Privacy, Proofs, and Compute may share that asset, but whether they should is an empirical architectural question rather than a premise.

The protocol and its price are different machines. The protocol machine supplies the triad and routes its demand through fees, burns, collateral, and constrained issuance. The market machine wraps the resulting asset into custody claims, funds, derivatives, leverage, and allocation rules. Either machine can move without the other. Price is something to explain, not evidence to cite.

Fees, burns, and collateral establish a competitively priced cash-flow claim and a balance-sheet floor—both of which a discounted-cash-flow valuation reproduces and neither of which is a monetary premium. The monetary claim rests on something else: a state-contingent **holder-side service flow**, the bearer’s ability to transact, prove, hold, and exit when the substitutes for proofs, courts, and custodians have stopped working. That mechanism is standard; this thesis argues it and declines to size it.

The competitor to an open verifiable stack is accordingly not fiat but a **competent closed sovereign stack**—state-integrated energy, industry, compute, payments, and identity—which may win on cost, uptime, and build speed while inverting who ends up sovereign. Two such stacks are now under construction rather than one. Duration-neutrality of the monetary object is a repression-resistance requirement, not a theory of capital formation: reserve collateral and the construction interval of plants are different functions, and proofs must not become the bond. The thesis publishes **fifteen red lines**: named, measurable conditions any one of which would retire it. [Section 7.1](#) states the argument as a chain of nine links and names the links it rates weakest.

Extended Abstract

Post-Bretton Woods money is increasingly backed by *compliance* and *regulatory enforcement* rather than reserves or convertibility. Debt stocks that cannot be honored in real terms make financial repression and balance-sheet capture arithmetically attractive, and the dependency now runs in both directions, as fiscal capacity itself comes to rest on the market value of the asset complexes the state regulates; synthetic media and platform moderation undermine the assumption that “seeing is believing”; AI concentrates cognition and compute behind hyper-scaler chokepoints; hardware, networks, identity, and custody systems are progressively integrated into state and platform surveillance infrastructure. In that environment, stores of value that depend on soft guarantees—custodians, editorial gatekeepers, platform labels, regulated wrappers, or vendor invoices—are brittle.

This thesis tests a narrower claim: **Privacy, Proofs, and Compute** can support an agency-preserving bearer asset that may earn monetary premium under repression, but only where the full stack delivers usable service under stress and a durable native holder base warehouses the asset’s risk. Fees, burns, collateral, disciplined issuance, and public telemetry establish value *accrual*—a competitively priced cash-flow claim and a balance-sheet floor, both of which a discounted-cash-flow valuation reproduces and neither of which is a monetary premium. The monetary claim rests on something else: a state-contingent **holder-side service flow**, the bearer’s ability to transact, prove, hold, and exit when the substitutes for proofs, courts, and custodians have stopped working. That mechanism is standard; this thesis argues it, declines to size it, and makes its physical, market, and holder conditions measurable so the question can be settled rather than asserted. We build a threat model that assumes intentional repression, not benevolence—including the softer **administrative repression** of custody defaults, compliance rules, and platform terms of service—and introduce a **seven-layer architecture** grounded in Layer 0 verifiable machines, energy, and physical capacity (Facility Capacity Receipts), extending through communications, software distribution, identity, truth & work, value & settlement, and a telemetry/governance layer that measures drift rather than denying it.

At the economic core, we formalize **verification asymmetry** and define **VerifyPrice** (the time-and-cost vector for independently checking a claim) as the hinge that turns proofs and verified FLOPs into commodities rather than platform IOUs. We propose a **Create/Compute Prove Settle Verify** loop, a modular stack of **twelve primitives** and **four reference applications** (private treasury & payroll, media provenance, verified inference, proof/compute procurement), and an operator/investor telemetry program (including extensions like **Veri-**

fyReach, VerifySettle, and the **Wrapper Dominance Ratio**) that makes neutrality, repression-resilience, and co-option risk falsifiable.

Even when these conditions hold at the protocol layer, observed asset prices can be dominated by external financial wrappers, price-insensitive allocations, mechanical rebalancing, and dealer hedging. This thesis therefore distinguishes **native value capture** from **market price capture** and introduces **VerifyFlow**, a public market-realization telemetry program measuring wrapper exposure, creations and redemptions, holder flow elasticity, leverage, concentration, dealer dependence, and the divergence between financial exposure and protocol-native use. Price is not treated as proof of monetary adoption. Recurring native fees, burns, collateral, settlement, and verified-work demand under healthy SLOs are what make value accrue to a triad asset at all, and are necessary before any monetary question can be asked of it; wrapper-led appreciation without those conditions is financialization rather than monetary validation.

Beneath both loops sits a physical dependency the thesis makes explicit rather than assuming away. Gross capacity is not deliverable service: the relevant quantity is the surviving flow through energy or fuel $\rightarrow$ firm electricity $\rightarrow$ transformer and switchgear $\rightarrow$ cooling and water $\rightarrow$ hardware $\rightarrow$ network $\rightarrow$ workload $\rightarrow$ proof $\rightarrow$ verification $\rightarrow$ settlement $\rightarrow$ usable service. We therefore report **Delivered Verified Capacity** beside **Sovereign Optionality**: the first measures scenario-specific surviving service flow and its minimum cut, while the second measures pathway diversity. Issuance claims are bounded by stress-adjusted deliverable capacity, not nameplate hardware or a benign-state average. Because verification affordability is denominated in real resources on reference hardware, it is exogenous to token price but *not* exogenous to the conditions under which power and unprivileged hardware can be obtained. The hinge is externally triggerable.

The **Topological Scarcity Lemma** follows: a nominal stock of resources, capital, or capacity does not secure a monetary service; the relevant quantity is the stress-deliverable flow through the narrowest non-substitutable edge connecting input to holder utility. The associated **Pressure–Capacity Corridor** writes the potential monetary service as $\Pi_{\text{monetary}}(R) \propto \Delta(R) \times A_{\text{full-stack}}(R) \times H(R)$, where differential demand may rise with regime pressure, full-stack availability may fall, and $H(R)$ is the holder’s practical ability to self-custody, transact, prove, and exit. The likely shape is an inverted U: little differential value under benign institutions, the strongest case while substitutes weaken but the open stack still works, and declining service when extreme pressure disables the stack or the holder.

The competitor to an open verifiable stack is accordingly not fiat alone but a **competent closed sovereign stack**—state-integrated energy, industry, compute, payments, and identity—which may win on cost, uptime, and build speed while inverting who ends up sovereign. Architecture is likewise not pre-selected: the thesis compares one native monetary asset, a neutral reserve plus service-specific credits, and shared settlement plus modular domain col-

lateral. Capability was never the claim; the properties are. In total the thesis publishes **fifteen red lines**: named, measurable conditions any one of which would retire it.

Gold, silver, Bitcoin, and physical hard assets remain load-bearing bridge assets while this stack is under construction; the triad becomes a candidate destination only when VerifyPrice, VerifyReach, VerifySettle, Layer 0 capacity telemetry, and value-capture metrics prove resilience under adversarial conditions. This is not a call to abandon hard assets prematurely, nor a claim that new money is inevitable. Bitcoin's work function is conceded to be monetarily superior to proof-of-useful-work, and for a reason that does not go away: nobody outside the system buys a hash, so nobody can subsidize, mandate, or withdraw demand for it, and a useful-work function necessarily introduces a buyer who can be. The claim made here is not that a better money is available, but that portable attestation and verified compute are separate goods with separate demand, and that whether they carry monetary premium or merely price as services is open. The result is not a single chain or guaranteed new money, but a **research and engineering agenda**: a "Bell Labs" for proof-of-useful-work and lawful privacy, testing whether Privacy, Proofs, and Compute can earn a store-of-value premium under measurable conditions—if verification remains cheap, privacy settlement remains usable, useful-work markets avoid capture, physical infrastructure remains legible, and protocol design converts recurring triad demand into scarce, non-bypassable, agency-preserving asset value.

Executive Memo: Why AfterFiat Exists

The fiat problem is no longer only inflation. It is the merger of debt arithmetic, compliance-backed money, platform-mediated truth, AI enclosure, and physical compute constraints.

Sovereigns with debt loads that cannot be honored in real terms will prefer repression to default. Platforms facing synthetic media and regulatory pressure will prefer managed truth to open verification. AI hyperscalers will prefer renting cognition to distributing it. Custodians and wrappers will prefer price exposure without native usage. Citizens below the participation line will be offered management rather than agency.

This thesis is the counter-architecture: private settlement instead of surveilled accounts, portable proofs instead of platform labels, verified compute instead of vendor promises, and public telemetry instead of trust. Capturing indispensable demand, and not being bypassed by fiat, stablecoins, cloud contracts, custodial wrappers, or administrative convenience, is what makes value *accrue* to the asset. It does not by itself make the asset money: a fee stream and a collateral lockup are quantities a discounted-cash-flow valuation reproduces. The monetary claim rests on a state-contingent holder-side flow—what the bearer can still do when the substitutes for proofs, courts, and custodians have stopped working—and this thesis argues that mechanism without sizing it.

The thesis is conditional, and two concessions belong at the front rather than in an objections chapter. Bitcoin’s work function is monetarily superior to proof-of-useful-work, because a puzzle nobody outside the system buys has no demand curve anyone can move. And gross native demand is not a monetary anchor: just-in-time fee purchases, burns, operator inventory, collateral requirements, and wrapper exposure do not establish a persistent, self-custodied, loss-bearing constituency willing to add through stress. Gold, Bitcoin, and hard assets are the bridge. Privacy, Proofs, and Compute supply different goods; whether one native asset, a neutral reserve plus service credits, or shared settlement with modular collateral best serves them remains an empirical question.

The protocol and its price are different machines.

The protocol machine supplies Privacy, Proofs, and Compute and routes their demand through fees, burns, collateral, and constrained issuance. The market machine wraps the resulting asset into custody claims, funds, derivatives, leverage, and systematic allocation rules.

The first machine determines whether the asset has monetary substance. The second determines how that substance—or the story about it—is represented in price.

Either machine can move without the other. A protocol can improve while its price falls during wrapper redemptions. Its price can rise while native use stagnates because an ETF,

treasury company, or leveraged product directs large exposure demand toward a thin market.

After Fiat therefore requires two forms of falsifiability: **stack telemetry** for monetary function, and **flow telemetry** for market realization.

A further unbundling belongs at the front rather than only in an objections chapter. Duration-neutrality of the monetary object is a repression-resistance requirement. It is not a theory of capital formation. Reserve collateral and long-lived project finance are different functions; Treasuries currently smash them together; proofs can audit the second and must not become it (Section 3.3, Section 31.5.1).

Both machines are bolted to a grid.

Cheap verification is the hinge of the entire argument, and it is priced in electricity and reference hardware. That makes it exogenous to token price but not to policy: a state that finds it awkward to ban verification can simply decline to energize it, and curtailment, rate discrimination, and interconnection queues survive judicial review comfortably. The exposures differ—verification is threatened by hardware obtainability, proving by bulk power—and conflating them produces overclaims in both directions. The thesis measures the substrate as **sovereign optionality**, prices its fragility into issuance, and retires itself if verification affordability becomes a sovereign policy variable rather than a market outcome.

The real competitor is competence, not fiat.

A state-integrated closed stack—sovereign energy, industry, compute, payments, and identity—is the same trust-minimization instinct applied to matter, and it may deliver better cost and uptime than anything open. The thesis does not contest that. It argues that non-custodial settlement, default privacy, portable identity, checkable receipts, and practical exit are what make an asset a store of value under repression, and that a closed stack withholds precisely those, because withholding them is what makes it closed. Identical dashboards, opposite answers on whether a user can leave. The greater risk is not losing to that stack but quietly becoming it, which is why agency and exit are on the scoreboard and why the thesis names **fifteen red lines** rather than promising success.

Why instrument rather than forecast.

Every prior restructuring of how civilization handles information arrived as an externality of something cheaper, and nobody ordered its consequences. The printing press was adopted to copy books. Nobody adopted it to create the silent private reader, standardized parts, or the nation-state; those arrived over centuries, as byproducts of a new habit at scale. The lesson generalizes, with the caveat that it is mechanism rather than measurement: a new medium is adopted for the function it makes cheaper, and it restructures whoever habituates to it, on a schedule nobody plans.

The stack tested here is being adopted to check work more cheaply. What cheap checking may do to the sensory basis of trust—to who can verify whom, and therefore to who must be trusted at all—is similarly unpredictable, and this thesis does not pretend otherwise. It re-

sponds by instrumenting the transition instead of forecasting it: fifteen red lines, public boards, and SLOs exogenous to token price, so that the consequences, good or bad, announce themselves in telemetry rather than in narrative ([Appendix J](#) records what each revision added). The frame is borrowed from the media-transition literature; the falsification conditions are not.

Contents

Abstract	ii
Extended Abstract	iii
Executive Memo: Why AfterFiat Exists	vi
Reader’s Map	xx
Thesis in Plain Language	xxv
The Conditional Chain at a Glance	xxvii
Key Definitions	xxx
I Context & Claims	1
1 Introduction	3
1.1 Central Claims	6
1.2 Premises	8
1.3 Contributions	9
1.4 End-to-End Vignette: The Loop in Action	11
2 The Failure of Soft Guarantees	15
2.0.1 A Prior Transition: What the Press Actually Changed	16
3 The World Forces New Monetary Primitives	18
3.1 Macro Playbook: Debt Repression Flight to Neutrality	18
3.1.1 The Treasury Market as Control Panel	19
3.1.2 Who Actually Buys the Paper	20
3.2 The Collateral Loop	23
3.2.1 The Circuit	24
3.2.2 Why This Explains Balance-Sheet Repression	24
3.2.3 A Branch in the Playbook	25
3.2.4 Why AI Sits at the Center of the Loop	25
3.3 Two Kinds of Duration	26

3.4	The Repression Playbook, Then and Now	29
3.5	Social: The Web's Trust Default Has Flipped	29
3.6	Political: The Participation Line Has Broken	31
4	First Principles: What a SoV Must Survive	32
5	Threat Model	36
5.1	Scope and Assets	36
5.2	Adversary Classes	37
5.3	Threats by Loop Stage	41
5.4	Mapping the Authoritarian Playbook	42
5.5	Energy & Physical Interdiction	44
5.6	Explicit Tensions and Boundary Conditions	46
6	Layers of the Cypherpunk Stack	50
6.1	Modules in Outline: Applications and Primitives	53
II	The Triad and the Monetary Candidate	55
7	The Triad and the Monetary Candidate	57
7.1	The Conditional Chain, End to End	58
7.2	Topological Scarcity and the Pressure–Capacity Corridor	62
7.3	The Triad Coherence Test	62
7.4	Monetary Objects and Value Capture	63
7.5	Three Reference Designs	67
8	Privacy as Private Money	69
9	Proofs as Attestation Money	72
10	Compute Through the “AI Money” Lens	74
10.1	Compute Deflation and What Remains Scarce	74
11	Work Credits: Energy-Anchored Claims	77
11.1	Definition: Work Receipts vs. Work Credits	77
11.2	Energy Anchoring	79
11.2.1	Layer 0 Maturity and Economic Consequences	79
11.3	Robots, AI, and the Demand for Work Money	80
11.4	Why Work Credits Remain Service Claims	81
11.5	Monetary Role	81

11.6 Issuance: Tying Credits to Real Capacity	82
11.6.1 The Supply Balance Equation	83
11.7 Energy-Priced, Not Energy-Pegged	84
11.8 Why Work Credits Are Typed Service Claims, Not Presumed Money	86
11.9 Economic Linkage: How Triad Demand Becomes Asset Value	87
11.9.1 Value Capture vs. Price Capture	91
11.10 Fee Incidence and the Level of the Fee	91
11.10.1 Incidence: A Fee Is a Wedge on Turnover, Not a Claim on Profit	92
11.10.2 The Base, and Why It Deflates	93
11.10.3 The Bound on the Sustainable Fee	94
11.10.4 Bypass Is a Magnitude, Not a Binary	95
11.10.5 What Determines Δ , and Why It Is Probably Thin	95
11.10.6 Two Bypass Channels the Thesis Has Not Been Counting	96
11.10.7 What the Five Conditions Establish, and What They Do Not	97
11.10.8 Where a Monetary Premium Would Actually Come From	98
11.10.9 Collateral: Unit-Elasticity, Floor, and Wrong-Way Risk	102
11.11 From Utility Demand to Monetary Premium	105
11.12 Why Users Cannot Simply Bypass the Asset	106
11.12.1 The Incumbent Bypass Channel Is Not Hypothetical	106
11.13 Wrapper Dominance Risk	108
11.14 The Market Realization Plane	110
11.14.1 A Three-Stage Model	110
11.14.2 The Numeraire Is Not Fixed	112
12 Service-Instrument Design Space	114
12.1 The “ZK Money” Analytical Lens	114
12.2 The “AI Money” Analytical Lens	115
12.3 Hybrid Instruments	115
13 The ZK + AI Service Economy	117
13.1 Why Open Hardware Is a Precondition for the ZK-Economy	117
13.2 How This Opens a ZK + AI Economy	118
14 SoV Evaluation Framework	120
14.0.1 Comparison: Traditional Hard Assets vs. the Base-Asset Candidate	122
III Infrastructure Layers 0–3	124
15 Layer 0: Verifiable Machines & Energy	126

15.1	Why Layer 0 Is a Monetary Question	126
15.2	Design Goals and Non-Goals	127
15.3	Hardware as Base Reality for Work Credits	128
15.4	The Layer 0 Feasibility Ladder	129
15.4.1	Policy Hooks: Economic Treatment by Grade	130
15.4.2	Quantitative Thresholds (Reference Design)	130
15.4.3	Pragmatic Starting Point (L0-A)	131
15.4.4	The Path Forward	131
15.4.5	Why This Prevents “Layer 0 Is Impossible, Therefore Thesis Fails” . . .	132
15.5	Concrete Components of Layer 0	132
15.5.1	Open Designs Where Possible	132
15.5.2	Attested Randomness and Entropy	132
15.5.3	Attestation Without Priesthood	133
15.5.4	Lot Sampling and Destructive Audits	133
15.5.5	How This Opens New Economies	135
15.6	Verifiable Power: Energy as First-Class Input	136
15.7	Facility Capacity Receipts: Beyond Energy	137
15.8	Sovereign Optionality: Aggregating FCR Into One Exposure	138
15.8.1	Adaptation of Scale	138
15.8.2	Mapping to Existing FCR Fields	139
15.8.3	Facility and Network Form	139
15.8.4	Governance of the Weights	140
15.8.5	What $\mathcal{O}_s$ Is For	140
15.9	Delivered Verified Capacity: Surviving Service Flow	140
15.10	Resilience Versus Affordability: The Layer 0 Cost Tension	143
15.10.1	Two Different Exposures	143
15.10.2	Disruption-Adjusted VerifyPrice	144
15.10.3	The Decision Rule	145
15.11	Operational Patterns: Profiles, Upgrades, and Failure Modes	147
15.11.1	Hardware Profiles and Workload Binding	147
15.11.2	Upgrades and Deprecations	147
15.11.3	PQ and Cryptographic Agility	148
15.12	Stress Tests for Layer 0	148
15.13	What Layer 0 Exports to Higher Layers	148
16	Layer 1: Reachability	150
16.1	Threat Model: What We Must Survive	150
16.2	Design Rules: Protocol Posture Under Pressure	151
16.3	Mechanisms: What Needs to Ship	151

16.4	How Layer 1 Anchors the Triad	151
16.5	VerifyReach: Communications Telemetry	152
16.5.1	VerifyReach Measurement Specification	153
17	Layer 2: Distribution & Execution	155
17.1	Threat Model & Objectives	155
17.2	Design Rules	156
17.3	Reference Distribution Architecture	156
17.4	Key Management & Release Process	157
17.5	Update Economics & Neutrality	157
17.6	Fallback Playbooks	158
17.7	Lawful Privacy in Distribution	158
17.8	Update Telemetry	158
18	Layer 3: Identity & Claims	159
18.1	Why Identity Must Decouple from Doxxing	159
18.2	Design Goals	159
18.2.1	Issuer Pluralism Without Issuer Anarchy	160
18.3	The Identity Kernel	161
18.4	Human Identity: Personhood and Compliance Without Dossiers	161
18.4.1	Sybil Resistance Menu	162
18.5	Machine Identity: Capability, Not Brand	163
18.6	Reputation: Priced Behavior, Not Personal Data	163
18.7	Patterns That Keep Identity Private and Usable	163
18.8	Failure Modes & Guardrails	164
18.9	How Identity Fits the Modular Stack	164
18.10	Identity/Reputation SLOs	164
18.11	Implementation Sketch	165
18.12	Part III Summary: Minimum Viable Stack	165
IV	Truth, Work, and Settlement: Layers 4–5	167
19	From Infrastructure to Economics	169
20	Layer 4: Truth & Work	170
20.1	What Layer 4 Is (and Isn't)	170
20.2	Verification Asymmetry Revisited	170
20.3	Canonical Workloads and Proof Types	171
20.3.1	Canonical Workload Definition Template	172

20.4 PoUW Design Patterns	174
20.5 Proof Factories and PaL	176
20.6 VerifyPrice Observatory	176
20.6.1 VerifyPrice Measurement Specification	177
20.7 VerifyPrice in Practice	179
20.7.1 The Constitutional SLO Band	179
20.7.2 Protocol Affordability SLOs (Operational)	180
20.7.3 Token-Quoted VerifyPrice (Market Signal, Not Target)	181
20.8 Physical VerifyPrice: Auditing the Infrastructure Behind a Workload	181
20.9 Layer-4 Stress Tests	182
20.10 Minimum Viable Layer-4 Economy	182
21 Layer 5: Value & Settlement	185
21.1 Settlement as a First-Class Workload	185
21.2 Design Constraints for Privacy Rails	186
21.3 The Privacy Rails Kit (PRK)	186
21.4 Settlement Safety: Refund and Bridge Invariants	187
21.5 VerifySettle: Settlement Telemetry	187
21.6 Layer-5 Stress Tests	188
21.7 Minimum Viable Layer-5 Corridor	188
22 The Modular Stack	190
22.1 Primitive Catalog: The Twelve Pieces	190
22.2 The Four Reference Applications at a Glance	193
22.3 Reference Application: Private Treasury & Payroll	194
22.4 Reference Application: Media Provenance	194
22.5 Reference Application: Verified Inference	195
22.6 Reference Application: Proof/Compute Procurement	195
22.7 Part IV Summary	196
V Governance, Telemetry & Neutrality	197
23 Layer 6: Governance & Telemetry	199
23.1 “No Dashboards, No Trust”: Public SLOs as Constitution	199
23.2 Control Surfaces: Parameters, Upgrades, Emergency Powers	200
23.3 Bell-Labs-Style R&D vs. On-Chain Governance vs. Off-Chain Norms	200
23.4 Governance as SLOs: The Five Invariants	201
23.5 Constitutional Enforcement	204
23.5.1 Machine-Enforced Constraints (Hard)	204

23.5.2	Override Path (Slow, Expensive)	204
23.5.3	Emergency Path (Narrow)	204
23.6	Monetary Constitution	205
23.6.1	Issuance Envelope	205
23.6.2	Fee Routing	206
23.6.3	Retirement Rule	206
23.6.4	Risk Haircuts	206
23.6.5	Coverage Target	208
23.6.6	Value Capture Loop	209
24	Extended Telemetry	211
24.1	Reference Verifier Classes	211
24.2	Reference Verifier Implementations	212
24.3	Cost Vector Definition	212
24.4	Sampling and Anti-Cherry-Pick	212
24.5	Data Availability and Anti-Memory-Hole	213
24.6	Reproducibility Contract	213
24.7	How to Measure: Receipts, Not Vibes	214
24.8	VerifyFlow: Instrumenting the Market Realization Plane	215
24.8.1	Wrapper-Level Vector	215
24.8.2	Asset-Level Vector	216
24.8.3	Holder Cohorts, Vintage, and Policy Concentration	218
24.8.4	Capital Survival: AUM Is Not Investor Return	218
24.8.5	Dealer Balance Sheets and Hedge Paths	219
24.8.6	Convenience-Yield Telemetry: Pricing the Holder-Side Service Flow	219
24.8.7	Buyer Quality and Price Provenance	221
24.8.8	Native Monetary Buyer Map	222
24.9	The Public Boards	223
24.10	Detecting Drift	225
24.11	Making Neutrality and Repression-Resilience Falsifiable	225
24.12	Policy-Attack Stress Harness	226
24.12.1	Market-Realization Stress Tests	226
25	Legal, Policy, and Jurisdictional Posture	228
25.1	Lawful Privacy as Protocol Design	228
25.1.1	What Lawful Privacy Is Not	228
25.1.2	Social Coercion Metrics: When “Optional” Becomes Mandatory	229
25.2	Securities Law: The Value-Capture Mechanism Is Its Own Legal Risk	230
25.3	Defense-in-Depth Against Bans and Sanctions	231

25.4 Labs, Foundations, and Neutral Router Commitments	232
26 Operator & Investor Checklist	233
26.1 Triad Supply: Privacy, Proofs, Compute	233
26.2 Value Capture & Anti-Bypass	233
26.3 Institutional Holdability	234
26.4 Stack Health: Layers 0–6	235
26.5 Telemetry Honesty	235
26.6 Red Flags and Failure Patterns	235
 VI Dynamics, Risk & Implementation	 237
27 Adoption Curve & Ecosystem Dynamics	239
27.1 The Transition Window: Hard Assets as Bridge	239
27.1.1 Portfolio Migration: From Bridge Assets to Triad Assets	240
27.2 Macro Policy State Model	240
27.3 Four Phases of Adoption	242
27.3.1 Phase Gates and Failure Gates	242
27.3.2 Adoption Curve To Expect	243
27.3.3 Phase I: Cypherpunk-Led (Prove the Plumbing in Public)	244
27.3.4 Phase II: Allocator-Led (Proofs and Privacy Become Budget Lines)	245
27.3.5 Phase III: Composability-Led (The Stack Disappears into Infrastructure)	247
27.3.6 Phase IV: Policy-Led and Path-Dependent (Verification as Public Good)	248
27.3.7 The Negative Case: Useful Infrastructure Without Money	249
27.4 Failure Gates	250
 28 Risk Analysis & Failure Modes	 251
28.1 Technical Risks	251
28.2 Economic Risks	251
28.3 Political Risks	252
28.4 Spec Drift, Vaporware, and Scoreboard Capture	253
28.4.1 Scoreboard Capture in the Market Realization Plane	253
28.5 Collateral Architecture Under Wrong-Way Risk	254
28.6 Red Lines: When the SoV Thesis Fails	254
28.7 Market Realization Warnings	263
 29 Implementation Sketches for Builders	 266
29.1 Minimum Viable Stack (MVS) Sequencing	266
29.2 Layer-0 Verifiable Machines	267

29.3 Privacy Rails: Making Settlement Safe and Boring	267
29.4 Proof Factories: Receipts as a Service	267
29.5 Compute Consensus Pilots	268
29.6 Developer Kit: Make “Import Proofs” the Default	269
30 The Closed Sovereign Stack	270
30.1 A State That Already Assumes Soft Guarantees Have Failed	270
30.2 Trust Minimization Applied to Matter	272
30.3 The Mirror: The Same Doctrine, Arriving in the West	273
30.4 Open Weights Are Not Open Sovereignty	275
30.5 The Inversion: Who Is Sovereign	277
30.6 What This Validates	278
30.7 What This Warns	278
30.8 What We Do Not Claim	279
30.9 What This Changes in the Stack	280
31 Objections & Responses	283
31.1 Monetary Objections	283
31.1.1 “Why not just Bitcoin?”	283
31.1.2 “This is just another chain / coin.”	285
31.1.3 “The triad does not belong in one token.”	286
31.1.4 “Utility does not imply money.”	287
31.1.5 “The asset can be bypassed.”	287
31.1.6 “Gold still beats this because gold needs no network.”	287
31.2 Technical and Infrastructure Objections	288
31.2.1 “Users don’t care about privacy or proofs.”	288
31.2.2 “PoUW will centralize / can’t compete with hyperscalers.”	288
31.2.3 “Proofs do not prove truth.”	289
31.2.4 “They’ll just shut off the internet / app stores.”	289
31.2.5 “This burns too much energy.”	289
31.2.6 “You will lose the competition for electrons.”	290
31.3 State, Political, and Governance Objections	291
31.3.1 “Governments will never allow lawful privacy at scale.”	291
31.3.2 “Receipts will become surveillance tools.”	291
31.3.3 “This just gives the state better compliance tools.”	292
31.3.4 “The closed sovereign stack simply wins.”	292
31.3.5 “A perfect ledger hardens an unequal starting state.”	293
31.3.6 “This is political preference dressed as analysis.”	294
31.3.7 “AI will be enclosed anyway.”	294

31.3.8 “Governance will just re-centralize.”	295
31.4 Market-Structure and Co-option Objections	295
31.4.1 “The asset will be co-opted, not bypassed.”	295
31.4.2 “This will just become another ETF.”	296
31.4.3 “The price went up, so the monetary thesis is working.”	296
31.4.4 “Wrappers increase access, so co-option is harmless.”	296
31.4.5 “Market mechanics are external, so protocol designers can ignore them.”	297
31.5 Completeness Objections	297
31.5.1 “Who finances the reactor?”	297
31.5.2 The Open Duration Warehouse	298
32 Why These Could Become Money	300
32.1 As a Store of Value	301
32.2 As Stack	302
32.3 As Telemetry	302
32.4 Trading One Base Reality for Another	303
33 Conclusion: A Bell Labs for Privacy, Proofs, and Compute	304
33.1 The Quiet Mic Drop	304
33.2 Final Thoughts	305
33.3 Two Loops, Not One	305
33.4 Equip, Don’t Manage	306
Sources	308
A Formal Model of Verification Asymmetry & VerifyPrice	317
A.1 Definition 1: VerifyPrice(W) and Verification Overhead	317
A.1.1 Implementation Note: VerifyPrice Telemetry Schema (Sketch)	319
A.2 Definition 2: Facility Energy Receipt (FER)	320
A.3 Definition 3: PIDL Receipt (Proof Interface Definition Language)	320
A.4 Definition 4: VerifyReach(N, R)	321
A.5 Definition 5: VerifySettle(C)	321
A.6 Definition 6: Repression Wedge	322
A.7 Definition 7: Liquidation Effect	322
A.8 Definition 8: Physical VerifyPrice	322
A.9 Definition 9: VerifyFlow	323
B Practical KPIs & Telemetry Templates	324
C The SDK (Proofs-as-a-Library)	326

D	Energy & Plant Architecture	329
E	Hardware Profiles	331
E.1	Concept	331
E.2	Naming	331
E.3	Hardware Profile Specification	332
E.4	Hardware Profiles in Receipts and Policies	332
E.5	Example Profiles	332
F	Communications Resilience Mechanisms	334
G	Glossary of Terms & Notation	335
H	Formal Model of Market Realization, Wrapper Flows, and Price Capture	342
H.1	Flow Elasticity	342
H.2	Leveraged Rebalancing and Flow-Induced Exposure	344
H.3	Net Mechanical Gain	344
H.4	The Recycling Boundary	345
H.5	Wrapper Recycling Ratio	345
H.6	Wrapper-Class Applicability	346
H.7	Return-Decoupled Allocation Flow	346
H.8	Market Impact and Liquidity	348
H.9	Volatility Drag on Leveraged Wrappers	349
H.10	Dealer Hedge Decomposition	349
H.11	Measurement Contract	350
H.12	Reference Stress-Harness Sketch	350
H.13	Boundaries and Caveats	350
H.14	Buyer Quality and Price Provenance	351
I	Scenario Analysis: The Collateral Loop Under Stress	352
I.1	Correlated Loss: A Recession With Rising Long Yields	352
I.2	The AI–Fiscal Collateral Loop	354
I.3	The National-Champion Conversion Pathway	354
I.4	Quarantined Claims	355
I.5	Indicators: What Would Confirm or Weaken This	355
J	What Changed Across Releases	359

Reader's Map

This thesis is structured as a ladder of seven interlocking layers. Each layer has one or more “home” sections, and most layers also surface in the **Create/Compute Prove Settle Verify** loop and in the telemetry program.

Use the following table as your compass:

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 0 Verifiable Machines & Energy	Physical base reality: open silicon, sampled supply chains, and the power plants that keep provers and routers alive. <i>For pure proof correctness, sound cryptography should tolerate untrusted provers. Layer 0 becomes essential when claims include physical capture, energy use, machine identity, witness confidentiality, or useful-work fairness.</i>	Chapter 15 <i>Layer 0</i> (open hardware, profiles, lot sampling, FERs); Chapter 29 (L0 implementation sketches); Appendix D (Energy & Plant); Appendix E (Hardware Profiles)	Hardware profiles (HIDs); lot-sampling coverage; Facility Energy Receipts (FERs); Facility Capacity Receipts (FCRs); sovereign optionality $\mathcal{O}_s$; Physical VerifyPrice (infra audit); disruption-adjusted VerifyPrice; PUE/ERE/WUE; kWh/receipt; profile incidents	RL10 (physical opacity) or RL13 ($\mathcal{O}_s^{\text{net}}$ below threshold / rationing jurisdictions); hinge VerifyPrice becomes sovereign-policy-driven

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 1 Reachability: Communications & Transport	Keeps packets flowing under DPI, filtering, and shutdowns. <i>Without Layer 1, the loop cannot run and “public verification” becomes theoretical.</i>	Chapter 16 <i>Layer 1</i> (threats, design rules, mechanisms, VerifyReach); Chapter 29 (comms sketch); Appendix F (Comms Resilience)	VerifyReach(N,R): p50/p95 time-to-first-connection; succ ₁ /succ ₂ reachability; failure causes; share encrypted/obfuscated traffic; peer-set diversity (geo/ASN); swap success under network stress	>20% ASNs unreachable for >7 days
Layer 2 Distribution & Execution: Software Supply & Runtime	Ensures honest clients and updates can be shipped even when app stores, CDNs, and DNS are weaponized. <i>Protects the code that speaks the protocol and runs PaL/PRK.</i>	Chapter 17 <i>Layer 2</i> (threats, design rules, multi-home architecture, update pipeline, key mgmt, fallback playbooks); Chapter 29 (distribution tooling); Appendix B (update KPIs)	Update health: p50/p95 metadata & artifact fetch times; success/rollback/signature-failure rates; mirror/CDN top-N share; geographic/ASN diversity; key-rotation & incident history	Single key compromise can push malicious updates

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 3 Identity & Claims: Humans & Machines Without Doxxing	Lets humans and machines prove capabilities and rights (age, residency, uniqueness, model ownership, device profile) without turning identity into a global dossier. <i>Reputation is receipts, not biographies.</i>	Chapter 18 <i>Layer 3</i> (identity kernel, human & machine identity, reputation, guardrails, how identity uses the loop); Chapter 25 (lawful-privacy requirements); Appendix C (PaL/PRK SDK)	Identity-predicate VERIFYPRICE (p50/p95 verify times for age/residency/set-membership proofs); unlinkability tests; reputation from PIDL receipts (SLA fulfillment vs slashes); hardware-anchored machine attestations	Identity requires centralized issuer
Layer 4 Truth & Work: Proof Systems and PoUW	Converts expensive work into cheap-to-check receipts. <i>This is where verification asymmetry lives and where VERIFYPRICE is defined.</i>	Conceptual: Chapter 10 <i>Compute Through the “AI Money” Lens</i> , Chapter 11 <i>Work Credits</i> ; Structural: Chapter 19 <i>Infrastructure to Economics</i> ; Chapter 20 <i>Layer 4</i> (canonical workloads, PoUW, proof factories, PaL); Chapter 22 (primitive catalog); Appendix A (formal VerifyPrice)	VERIFYPRICE(W) per workload: p50/p95 verify times/costs; failure rates; $r(W) = v(W)/p(W)$; prover concentration; SLA attainment (Bronze/Silver/-Gold); MatMul-PoUW economics; Work-Credit issuance	VerifyPrice p95 exceeds SLO for ≥ 3 months

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 5 Value & Settlement: Privacy Rails & Non- Custodial Flow	Moves value non-custodially and privately, with auditability via receipts. <i>This is where the service markets described by the legacy “Private Money,” “AI Money,” and “Proof/Attestation Money” lenses settle; their Work Credits and capacity instruments remain typed service claims. The separate base asset is the only conditional monetary candidate.</i>	Analytical lenses: Chapter 8 Privacy as Private Money , Chapter 10 AI Money , Chapter 9 Proof/Attestation Money ; service claims: Chapter 11 (Work Credits) ; Stack: Chapter 21 Layer 5 (settlement as workload, PRK, bridge-safety, corridors); Chapter 22 (treasury/payroll, media, inference, procurement flows); Chapter 29 (privacy toolkit); Appendix F (swaps, routing)	VERIFYSETTLE(C) : swap success $\geq$ target; refund_safe(C) = 1.0 ; p50/p95 time-to-finality; anonymity-set size & churn; corridor LP/route concentration; non-custodial vs custodial flow share	Corridor refund_safe < 100% (repeated)

Continued on next page...

Layer	Role in the Stack	Primary Sections	Key Metrics / Concepts	Falsification Test
Layer 6 Governance & Telemetry	The immune system and constitution: keeps drift and capture visible and forces responses. <i>Turns “trustlessness” into dashboards and runbooks.</i>	Metrics: Chapter 20 (VerifyPrice observatory); Chapter 24 <i>Extended Telemetry</i> ; Governance: Chapter 23 <i>Layer 6</i> (SLOs as constitution, control surfaces, Bell-Labs vs on-chain vs norms); Legal: Chapter 25 ; Ops: Chapter 26 (checklist & red flags); Dynamics: Chapter 27 and Chapter 28 ; Appendix A and Appendix B	VerifyPrice dashboards per workload; VerifyReach & VerifySettle boards; VerifyFlow (market realization); twelve public boards or maps (Section 24.9), including native-buyer, common-cause, and sponsor-dependency telemetry; decentralization metrics; fee+burn coverage; Wrapper Dominance; agency preservation	Receipt datasets unverifiable or controlled by single party; wrapper-led price without native use (interpret via VerifyFlow, not as sole red line)

Readers who want the version history—what each release from v1.2 onward added, and why—will find it in [Appendix J](#). It is not needed to read the argument.

Thesis in Plain Language

The core claim: Under sustained administrative and financial repression, a bearer base asset may earn monetary premium if holding it preserves private settlement, portable proof, and access to verified compute after ordinary substitutes weaken. The full service path must remain stress-deliverable; demand must not bypass the asset; a persistent, self-custodied, loss-bearing holder constituency must absorb residual risk; and infrastructure credit must remain separate. Privacy, Proofs, and Compute may share the asset, but the thesis does not assume they should.

The accrual mechanism: Service value can accrue to the base asset through four channels:

1. fees paid in the native asset for proof generation, privacy settlement, and verified compute;
2. burns that permanently remove supply when capacity is consumed;
3. required collateral for provers, routers, and liquidity providers; and
4. scarcity constraints tied to energy and hardware, not fiat decree.

The value-capture bridge: Utility demand alone does not create store-of-value premium. Five conditions are required even to route service value to the asset:

1. the asset is the **required fee medium** for core triad services;
2. a meaningful share of fees is **burned or retired**;
3. operators must **stake the asset as collateral**;
4. issuance is **capped or capacity-constrained**; and
5. users **cannot obtain equivalent service quality** through bypass channels that avoid the asset.

If any condition fails, the system may be useful infrastructure while the base asset fails even as a service-value claim. If all five hold, the result is a cash-flow claim and collateral floor, not yet money.

The monetary bridge adds two requirements the accrual bridge cannot supply: stress-deliverable service and a holder constituency that retains self-custodied balances, bears losses, and accumulates through relevant pressure rather than buying fees just in time or holding wrappers.

The falsifiable test: It runs one way. If VERIFYPRICE rises, VERIFYREACH falls, VERIFYSETTLE breaks, DVC collapses, the capturable wedge closes, the native holder base proves procyclical, or the collateral-capacity spiral breaches Red Line 15, monetary candidacy fails. Healthy readings merely keep the candidate alive; no accumulation of green dashboards proves monetary premium.

What price does not tell you: None of the tests above is a price test. An asset's price can rise because ETFs, treasury companies, leveraged products, passive mandates, or dealer

hedging direct large exposure demand at a thin market, while native fees, burns, settlement, and proof purchases stay flat—and it can fall while all of those improve. That is why the thesis adds a fourth measurement family, **VerifyFlow** (Section 24.8), to separate protocol-native demand from wrapper-led exposure. Price is something to explain, not evidence to cite.

The thesis examines the triad from three complementary angles, plus one plane around them:

- **Angle 1 Monetary:** whether a bearer base asset can earn a next-generation store-of-value premium by preserving agency through Privacy, Proofs, and Compute; the service claims themselves remain non-monetary.
- **Angle 2 Stack:** the seven-layer cypherpunk stack that actually supplies these capacities.
- **Angle 3 Telemetry & Governance:** VerifyPrice/Reach/Settle + ops as the thing that keeps “repression-resilient neutrality” falsifiable.
- **The Market Realization Plane:** the external financial machinery—wrappers, custody, leverage, dealers, allocation rules—that determines how, when, and with what volatility any of the above shows up in price (Section 11.14).

The Conditional Chain at a Glance

The thesis establishes its claim through a conditional chain rather than an assertion. The chain is defended link by link across six Parts, which makes it easy to lose. This page states it once; [Section 7.1](#) walks it in full, says where each link is argued, and closes with [Table 7.1](#), which rates the argument link by link.

Utility demand → standardized work → receipts → stress-deliverable service → markets → fee flows → a scarce asset (only under non-bypassable accrual) → a persistent, self-custodied, loss-bearing, regime-responsive holder constituency → a possible store-of-value premium.

Every arrow is a conditional. None is asserted as inevitable, and the chain is only as strong as its weakest link. Each link therefore comes with the condition that would sever it:

	Link	Severed if
1	Utility demand exists and persists, through every cycle and especially under repression	the demand turns out to be cyclical discretionary spend rather than structural necessity
2	Demand standardizes into canonical work—workload registries, hardware profiles, SLA tiers	workloads stay bespoke, so no fungible unit forms and there is nothing to price
3	Work produces cheap-to-check receipts, with $r(W) = v(W)/p(W) \ll 1$	verification cost creeps toward production cost, and “anyone can verify” collapses into “trust the prover” (Red Line 1, the hinge)
4	Receipts represent stress-deliverable service, measured by scenario max-flow / min-cut across the full stack	nameplate capacity survives on paper while substitution latency or common-cause failure collapses usable service
5	Receipts enable markets: proof factories, neutral routers, SLA escrow, capacity procurement	routers, provers, or matching engines capture the market and rent replaces work (Red Line 3 territory)
6	Markets produce recurring fee flows, visible as fee coverage against the security budget	fee coverage collapses and workload demand proves speculative rather than budgeted
7	Fee flows accrue to a scarce asset— <i>only</i> under the five non-bypassable conditions of the Value Capture Lemma	users obtain equivalent triad capacity through hyperscalers, stablecoins, or custodial APIs without touching the asset (Red Line 6)
8	A persistent, self-custodied, loss-bearing, regime-responsive holder constituency anchors the asset	native demand is just-in-time, immediately sold, wrapper-led, leveraged, or procyclical
9	A scarce, captured, anchored asset may earn a store-of-value premium, if it stays liquid, neutral, verifiable, legally holdable, and agency-preserving	any store-of-value requirement fails, or the telemetry showing it is captured (Red Line 4)

Where it is weakest. Standardized work and proof/compute market formation remain engineering risks. Three bridges are explicitly *Medium*: stress-deliverable service, because common-cause and substitution data are immature; non-bypassability, because the economics requires a magnitude rather than a binary; and the holder anchor, because native transactions do not prove risk absorption. The store-of-value premium derives from link 9 plus state-contingency, not from fee accrual ([Section 11.10](#), [Section 11.10.8](#)). The full ratings are in [Table 7.1](#).

What the chain does not claim. Completing all nine links would establish only that the base asset is a defensible monetary candidate. It would not establish when, or with what volatility, price reflects that ([Corollary 11.1](#)). Nor is the chain suspended in mathematics: the service links consume bulk power and hardware, and the holder link requires balance sheets capable

of bearing loss. Both can fail while gross capacity or gross native demand still looks abundant.

Key Definitions

Before proceeding, we establish precise definitions for the core constructs that recur throughout this thesis. These are not metaphors; they are operationally specified primitives.

Definition: VerifyPrice(W) Specification Stub

For a canonical workload W , **VerifyPrice** is the public KPI vector:

$$\text{VERIFYPRICE}(W) \equiv (p_{50,t}(W), p_{95,t}(W), p_{50,c}(W), p_{95,c}(W), \text{fail}(W))$$

Where:

- $p_{50,t}(W), p_{95,t}(W)$: median and 95th-percentile verification **time** (seconds)
- $p_{50,c}(W), p_{95,c}(W)$: median and 95th-percentile verification **cost** (see cost vector below)
- $\text{fail}(W)$: verification failure rate (fraction of attempts that fail or timeout)

This is the canonical five-field tuple; the cost fields derive from per-run receipts via the published price model of [Section 20.6](#) (per-run cost, then percentiles — never recomputed from percentile inputs), and the physical-component tuples published alongside are views of this vector, not competing definitions ([Appendix A](#)).

Why this matters: VerifyPrice is the hinge that determines whether proofs and verified compute behave as commodities (publicly checkable) or as platform IOUs (trust someone’s claim). If $r(W) = v(W)/p(W) \ll 1$, verification is cheap relative to production and markets can form; if $r(W) \rightarrow 1$, we’re back to “trust the prover.”

Definition: Work Credits

A **Work Credit** is an energy-anchored claim on a standardized unit of triad work (privacy settlement, proof generation, or verified compute) that has been produced and attested under public SLOs.

Issuance: Credits are minted only when:

1. A valid proof of workload W at tier T is accepted by the network.
2. Telemetry confirms $\text{VerifyPrice}(W, T)$ and other SLOs (latency, failure rate, decentralization) are within bounds.

Claim semantics: Implementation-dependent. Work Credits can be designed across a service-contract spectrum:

- **Non-redeemable evidence-linked claims:** transferable claims referencing historical work. Scarcity may support price, but does not make them money.
- **Redeemable vouchers:** credits burnable for future proofs, compute, or settlement capacity. Provides direct utility claim.
- **Fee/collateral/governance medium:** credits required for network operations:
 - *Fee prepayment:* credit burns in lieu of per-call fees.
 - *Collateral:* credit staked as skin-in-the-game for provers, routers, and LPs.
 - *Governance weight:* credit-weighted voting in telemetry disputes and parameter changes.

These options are not mutually exclusive; a single network may support multiple redemption paths for different use cases.

Energy anchoring: Marginal cost of minting one credit is bounded below by energy and hardware required to pass verification. The difference from SHA-256 PoW is that this work has an *external buyer*—which is the point of the design and also its cost, since a buyer can be regulated, subsidized, or coerced in a way that a hash puzzle’s absent buyer cannot (Section 31.1.1). Each credit references a Facility Energy Receipt (FER) chain; if the referenced plant drifts out of profile ($PUE > 1.5$, carbon intensity $>$ threshold, etc.), downstream credits are flagged.

Deliverability bound: Issuance is capped by stress-adjusted DVC for the specified workload and tier, not by gross energy, installed hardware, or benign-state proof throughput.

Non-debt property: Work Credits do not promise fixed coupons or redemption in fiat terms. Value floats with demand for triad capacity.

Failure mode: If VerifyPrice regresses materially, new issuance halts until SLOs recover. Existing credits remain valid but may trade at a discount, reflecting the network’s degraded utility.

Definition: Instrument Hierarchy

The thesis distinguishes five categories of objects:

Evidence Objects (not scarce, not money):

- **PIDL Receipt:** Proof that a specific interaction occurred. Copyable, verifiable, not scarce.
- **Work Receipt:** Proof that a unit of work was completed. Evidence of past work; confers no future rights.
- **FER/FCR:** Evidence about energy, infrastructure, resilience, and the edge capacities used to compute DVC.

Capacity and Service Objects (typed claims, not presumed SoV):

- **Work Credit / WC-Voucher:** Prepaid or transferable access to specified triad capacity under workload, location, hardware, SLA, and delivery terms. May expire. Useful for service procurement and hedging, not presumed savings.

Duration-Bearing Credit:

- **Project Notes and Capacity Bonds:** Explicit credit claims with maturity, covenants, default states, and loss waterfalls. Proof-audited, but never transferred onto the base asset.

Derivatives and Operating Claims:

- **LP/Staking Shares:** Positions granting fee, slashing, corridor, or validator exposure. They are not monetary objects.

Conditional Monetary Candidate:

- **Base Asset:** The native fee and settlement unit. It is evaluated for monetary premium only if DVC, non-bypassability, holder quality, agency, and the remaining chain conditions pass.

Hierarchy rule: Throughout this thesis, “the asset” refers to the **base asset** unless otherwise specified. Evidence, service claims, duration-bearing credit, and derivatives do not inherit its conditional monetary candidacy.

Definition: Lawful Privacy

Lawful privacy is the design principle: *default privacy with optional, user-controlled disclosure*.

Concretely:

- **Default state:** Transactions, identities, and flows are encrypted and unlinkable without explicit consent.
- **Disclosure mechanisms:** Viewing keys, auditable receipts, and selective-disclosure proofs allow holders to prove specific facts (e.g., “I paid X to Y for purpose Z”) without exposing the full transaction graph.
- **No backdoors:** The protocol has no master keys, regulatory escrow, or “lawful intercept” APIs. Disclosure is always at the holder’s discretion.

Coercion boundary: Lawful privacy is a *technical* guarantee. It cannot prevent social or legal coercion to disclose viewing keys. What it guarantees is that (1) non-custodial routes exist, (2) disclosure cannot be forced at the protocol level, and (3) coercion surface is minimized by keeping data encrypted by default.

Quick Reference: Political Economy & Physical Capacity

These terms carry the thesis's political-economy and physical-capacity vocabulary. Full treatments appear at the cited locations; short forms are collected here for quick reference.

- **Balance-Sheet Repression** (Section 3.1.1): financial repression implemented through collateral rules, capital treatment, stablecoin reserve rules, custody mandates, and institutional balance-sheet incentives.
- **Administrative Repression** (Section 5.2): the conversion of formally optional financial, identity, compute, and settlement rails into practically mandatory rails through custody defaults, compliance rules, app-store control, tax treatment, institutional mandates, benefit systems, and platform terms of service.
- **Participation Line** (Section 3.6): the household or organizational threshold below which a person or firm lacks the redundancy to act freely across time—to fail, retry, transact, move, learn, refuse coercive terms, or survive shocks.
- **Agency-Preserving Infrastructure** (Section 3.6, Item 9): infrastructure that expands a user's capacity to act without converting the user into a dossier, dependency object, or platform account.
- **Homestead Ratio** (Section 5.6): the share of verified compute, proof generation, and AI-service capacity supplied by open-admission, non-hyperscaler, geographically diverse, independently verifiable operators.
- **Facility Capacity Receipt (FCR)** (Section 15.7): a signed, auditable receipt that extends Facility Energy Receipts with grid, cooling, redundancy, hardware, jurisdiction, and infrastructure-resilience claims.
- **Physical VerifyPrice** (Section 20.8, Appendix A.8): the time, cost, and confidence required to verify the physical infrastructure claims behind a unit of verified work.
- **Wrapper Dominance Ratio (WDR)** (Section 11.13): a reported *pair* — custodial or synthetic exposure share (stock) and wrapper-to-native activity ratio (flow) — never divided into one scalar. A rising WDR indicates the asset may be financializing faster than it is becoming money.

Quick Reference: Physical Sovereignty & Substrate

These terms make the physical layer falsifiable and name the closed-stack competitor. Full treatments appear at the cited locations.

- **Energy & Physical Interdiction** (Section 5.5): adversary class using curtailment, rationing, tariff discrimination, interconnection denial, and load prioritization to raise verification cost without prohibiting cryptography.

- **Sovereign Optionality ($\mathcal{O}_s$)** (Section 15.8): capacity-weighted index from Facility Capacity Receipt fields; feeds risk haircuts and Red Line 13.
- **Disruption-Adjusted VerifyPrice** (Section 15.10.2): probability-weighted verification cost across physical disruption states; resilience is subordinate to constitutional VerifyPrice SLOs.
- **Red Line 13: Energy Sovereignty Failure** (Section 28.6): verification affordability becomes a sovereign policy variable; monitors upstream of Red Line 1.
- **Closed vs. open sovereign stack** (Chapter 30): convergent trust-minimization applied to matter; inverted locus of sovereignty (state vs. participant).
- **Physical VerifyPrice (two senses)** (Appendix A.8, glossary): constitutional real-resource SLO vs. cost to audit FCR infrastructure claims—Red Lines 1 and 10 respectively.

Quick Reference: Market Realization & Price Formation

This is the market-structure vocabulary for reasoning about price without confusing it with adoption. Full treatments appear at the cited locations.

- **Native instrument vs. external financial wrapper** (Section 7.4): native instruments include the base asset, Work Credits, evidence artifacts, project credit, staking, LP, and corridor claims. Touching the protocol does not make an instrument monetary. External financial wrappers (spot ETFs, ETPs, treasury companies, custodial balances, futures, options, swaps, leveraged ETPs) do not touch it. Unqualified “wrapper” means the latter.
- **Value capture vs. price capture** (Section 11.9.1, Corollary 11.1): value capture routes native demand through the monetary object; price capture is demand for exposure to the asset’s price and can occur with no protocol use. Neither implies the other, in either direction.
- **Market Realization Plane** (Section 11.14): the external financial machinery through which claims on the native asset are represented, financed, allocated, and priced. Orthogonal to the seven-layer stack—a plane around it, never a Layer 7.
- **VerifyFlow** (Section 24.8, Appendix A.9): the fourth verification family, measuring the external financial representation and price-transmission state of the native monetary object.
- **Flow elasticity (ε)** (Appendix H.1): the responsiveness of a wrapper’s shares outstanding to changes in its value per share; the observable signature of holder operating rules. Measured, not structural.
- **Net mechanical gain (κ)** (Appendix H.3): $L(L-1) + \varepsilon L^2$, the coefficient converting an

underlying return into a wrapper's return-coupled exposure demand. Requires $L > 1$ (levered); the unlevered case is below.

- **Wrapper Recycling Ratio (WRR)** (Appendix H.5): $-\varepsilon L / (L - 1)$; equal to 1 when holder flows exactly offset gross rebalancing, below 1 when the wrapper amplifies, above 1 when it is countercyclical. Defined for levered wrappers only ($L > 1$): for unlevered wrappers ($L = 1$) the denominator vanishes, the wrapper has no rebalancing trade to recycle, and the published object is ε itself with the label “unlevered” (Appendix H.5).
- **Return-coupled vs. return-decoupled flow** (Appendix H.7): return-coupled flow responds to the return and shapes volatility and persistence; return-decoupled flow arrives independent of it and shapes destination and level. The decomposition is exhaustive and exclusive: the ε -component belongs to the coupled side (inside κ), the intercept α to the decoupled side, and the residual is published unabsorbed (Appendix H).
- **Wrapper-Native Growth Gap (WNG)** (Section 24.8): $\Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$, computed on unit-denominated quantities over a trailing four-quarter window (revaluation stripped); persistently positive means financial exposure is outgrowing monetary use.
- **Compositional adversary** (Section 5.2): a failure mode with no malicious actor, in which locally rational operating rules combine into concentration, procyclicality, recursive leverage, and uninformative prices.
- **Capital Survival Ratio (CSR)** (Section 24.8.4): whether successive investor cohorts in a wrapper preserved capital, as distinct from whether the product preserved AUM.
- **Policy Concentration Ratio** (Section 24.8.3): the share of net asset demand governed by the largest common rebalancing templates, treasury algorithms, or agent policies.
- **Automatic vs. marginal buyer** (Section 3.1.2): formula demand versus residual price-setting demand; quantity of bids is not quality of risk absorption.
- **Duration of the claim vs. duration of the project** (Section 3.3): a repression-resistant store of value must not be a coupon; civilization still needs a duration warehouse. Not a layer.
- **Buyer-quality checklist** (Section 24.8.7, Appendix H.14): persistence, horizon, liability match, loss-bearing capacity, countercyclicality—fields, not a scalar.

Part I

Context & Claims

This Part sets the stage: it explains why soft guarantees (central banks, broadcast media, credentialed authority) are failing under debt, AI, and platform control, and why we need new monetary primitives. It introduces the triad (Privacy, Proofs, Compute), the threat model, and the layered “cypherpunk stack” that the rest of the thesis will fill in.

Chapter 1

Introduction

Start Here: Four Reading Paths

This chapter opens with the cultural and historical argument for why anchors are being sought at all. That framing matters, but it is not the load-bearing part of the thesis, and readers who already accept the premise should skip ahead rather than wade through it.

If you allocate capital or underwrite risk: read the *Executive Memo* and *Reader's Map* in the front matter, then [Section 1.2](#) (the ten premises), [Section 7.1](#) (the argument in one page, with its weakest links named), [Lemma 11.1](#) (the condition that decides whether any of this captures service value), [Section 28.6](#) (the fifteen ways to prove us wrong), and [Chapter 14](#) (the evaluation questions to put to any candidate instrument).

If you build infrastructure: start at [Chapter 6](#) for the seven layers, then [Section 22.1](#) for the twelve primitives and [Chapter 22](#) for the four reference applications. [Chapter 29](#) is the 12–36 month build sketch.

If you are here to find the weak points: [Section 5.6](#) states the boundary conditions we cannot engineer away, [Chapter 31](#) answers the strongest objections we know of—including the completeness objection in [Section 31.5.1](#)—[Chapter 30](#) concedes what a competent closed sovereign stack does better, and [Table 7.1](#) rates our own argument link by link.

If you want the whole claim in one sentence: *Thesis in Plain Language*, in the front matter.

One warning that applies to every path: nothing in this document treats price as evidence. The monetary claim is tested on protocol telemetry, and price is interpreted separately ([Corollary 11.1](#), [Section 24.8](#)).

Every monetary epoch begins with an argument about what is real. James Dale Davidson and William Rees-Mogg, in their seminal work *The Sovereign Individual*, argued that the decentralization of computer networks would erode centralized power, and that governments, in their death throes, would reach for more aggressive tools of control: *capital controls, nationalization, outright authoritarianism*.

The old guarantees (central banks, broadcast media, credentialed authority) no longer hold

their shape under the pressure of digital networks and foundation models. They are under attack, not by a single conspirator, but by the physics of decentralization itself, which dissolves the monopolies that once defined our shared reality.

Marshall McLuhan saw an even deeper shift: the move from print, a one-to-many medium that centralizes narrative, to electronic networks, a many-to-many mesh that fragments it. A print-created reality is a centralized reality: he who controls the press, controls the narrative. This underwrote 20th-century monetary supremacy and global settlement currency backed by a monopoly on violence: the post-1971 US petrodollar.

In the 20th century, citizens got their singular cultural feed from *The New York Times* over breakfast, and later their dualistic, pseudo-antagonistic programming from CNN and Fox News. In the 21st, the internet (many-to-many) erodes this centralization of narrative. Take the network formerly known as Twitter: authoritarian regimes that attempt to censor information find it very hard to stop leakage. Within minutes, people with “smart” phones can produce a preponderance of evidence that either supports or invalidates the prevailing story.

But a countervailing force has arrived: AI. The quality of realistic generated content has crossed a critical threshold, approaching the asymptote of believability. Synthetic faces, voices, and scenes now compete with direct sensory experience.

This creates a world wealthy in symbols but poor in anchors: a sea of abstractions with no obvious way back to sensible shores.

So we look for new anchors.

Cypherpunk cryptography is the unifying answer: privacy by default, proof by construction, and compute that only gets paid when anyone can verify it. It replaces authority with protocols and gossip with receipts. In this frame, money is not a promise from a platform; it is what remains when verification is cheap and permission is irrelevant. Cypherpunks don’t ask for integrity; they **instrument** it. They don’t trust platforms; they price receipts. The triad is that credo made economic: privacy that preserves agency, proofs that travel, and compute that earns only when anyone can verify it.

Bitcoin’s SHA-256 puzzle mints digital scarcity by tying consensus to thermodynamic cost, and it is worth being exact about why that works: the work has *no external buyer*. Nobody outside the system wants a partial hash preimage, so nobody can subsidize it, mandate it, price it, or withdraw demand for it. The cost is therefore non-negotiable, and the money inherits that objectivity. The design direction explored in this thesis keeps what makes PoW legitimate (open admission and unpredictable leader election) but swaps the work function, so that the “lottery ticket” is earned by producing succinct, publicly verifiable receipts of compute somebody wanted. Miners win by attaching a proof that a market-demanded computation was done correctly.

A natural anchor workload is matrix multiplication. MatMul-PoUW constructions make

verification asymptotically cheaper than naive production (for example, $O(n^2)$ verification vs. $O(n^3)$ multiplication) while keeping verifier overhead at $(1 + o(1))$ relative to the best-known randomized checker. That turns AI’s core primitive into **verified FLOPs**: a commodity unit anyone can check cheaply. Projects like Nockchain show the zk-PoW design space operating in the wild (fair-launch ethos, scope-minimized “dumbnet,” explicit issuance), illustrating that proof-carrying work can be wired into consensus without appointing gatekeepers. Keep the lottery; change the work. Make the prize a receipt the public can verify—and accept, in exchange, that the receipt has a buyer and the buyer has a jurisdiction.

Caveat: MatMul-PoUW constructions and useful-work consensus are promising research directions, not yet proven production monetary bases. The arXiv literature describes the optimal-security claim as a *conjecture*; deployment-grade engineering remains ahead. This thesis frames PoUW as the destination of a research program, not a shipping product.

Operationally, the goal is simple: bind useful work to the hash race **without introducing trust**.

Two deployable patterns are in scope:

- **(A) Hash-gated useful work.** A SHA-family threshold confers short-lived eligibility; a block is valid only if it carries a PoUW artifact (e.g., MatMul proof or zk-proof) seeded by header randomness.
- **(B) Proof-first selection.** Miners race to post useful-work receipts to a mempool; header entropy resolves ties and timing.

Both patterns must prevent precomputation (epoch seeds, commit-reveal), avoid closed-hardware dependencies, and expose verifier-light clients plus dispute/slash routes for junk artifacts.

When blocks routinely carry proofs that clear public SLOs and decentralization telemetry remains healthy (time-to-first-proof, top-N share, geo/ASN spread), block rewards bind issuance to capacity the world already buys: ZK proving, matrix multiplication, verified inference. So **Privacy, Proofs, and Compute may begin to earn monetary premium**—if the conditions outlined in this thesis hold.

State the cost of this at the outset, because it is real and it runs the other way. An external buyer of the work is a party who can be regulated, subsidized, or coerced, and whose demand can therefore be moved by something other than the resource cost of the work. That is precisely the exposure SHA-256 avoids by having no buyer at all, and it is why Bitcoin’s work function is the monetarily stronger construction. Proof-of-useful-work is not an improvement on it; it is a different trade—**monetary objectivity exchanged for capacity relevance**—made in the belief that the second is what a dense digital civilization will be short of. The trade is argued at full strength, against itself, in [Section 31.1.1](#), and a reader who wants the concession before the argument should go there first.

The rest of this Part translates that intuition into claims, contributions, a threat model, and a layered architecture.

What Would Falsify This Thesis?

The store-of-value claim fails if any of the following persist without remediation: (1) verification becomes expensive or gated (VerifyPrice breaks); (2) refund safety fails on admissible corridors; (3) verification concentrates in a single hardware profile or jurisdiction; (4) telemetry is captured or becomes unverifiable; (5) fee coverage collapses and workload demand becomes purely speculative; (6) users can consume triad capacity without touching the native asset; (7) wrapper and custodial exposure grows while native usage stagnates; (8) physical infrastructure claims become unverifiable; (9) AI compute enclosure persists; (10) lawful privacy collapses socially; (11) verification affordability becomes a sovereign policy variable; (12) the capturable wedge falls below the protocol take; or (13) a native-collateral drawdown drives operator exit, declining Delivered Verified Capacity, released float, and unrecovered service SLOs. These are illustrative; the full set of fifteen red lines is developed in [Chapter 28](#).

1.1 Central Claims

In this thesis we argue:

1. **Post-Bretton Woods money relies increasingly on compliance infrastructure.**

As convertibility receded, taxation capacity, legal tender status, central-bank balance sheets, collateral rules, and the compliance perimeter became the practical support structure of fiat money. When compliance becomes weaponized and asset freezing becomes policy, neutral stores of value become necessary infrastructure, not ideological luxuries.

2. **Three cryptographic capacities may support a monetary candidate under measurable conditions.**

Privacy (censorship-resistant settlement that preserves agency), **Proofs** (portable attestations of computation and provenance), and **Compute** (useful work wrapped in succinct guarantees) are distinct services. They support a monetary candidate only if their service path remains stress-deliverable, their demand cannot bypass the candidate asset, and native holders warehouse its residual risk. They need not share one token.

3. **A modular stack can operationalize this triad.**

We propose four reference applications (private treasury & payroll, media provenance, verified inference, proof/compute procurement) built on twelve primitives, with **verification asymmetry** and **VerifyPrice** as the key economic metrics that turn proofs and verified FLOPs into commodities rather than platform IOUs.

4. **Verifiable machines (Layer 0) are essential for physical-world claims.**

For pure proof correctness, sound cryptography should tolerate untrusted provers. Layer 0 becomes essential when the claim includes facts about physical capture, energy use, machine identity, witness confidentiality, side-channel resistance, or useful-work fairness. Open or sampled hardware is not ornament; it is the base layer that makes physical-world claims credible.

In the rest of the thesis we make this concrete in modular form. Four reference applications exercise the stack (private treasury & payroll, media provenance & authenticity, verified inference as AI service, and proof/compute procurement), supported by a small toolkit of primitives: a proofs-as-a-library SDK (“PAL”) that compiles claims to proofs, a privacy-rails kit that executes non-custodial, refund-safe settlement over BTC↔ZEC/XMR corridors, a minimal receipt schema (“PIDL”) that turns every proof or settlement into a portable artifact, neutral router logic that keeps useful-work markets open, and a VerifyPrice observatory that measures how cheap verification really is. Later sections develop these primitives and applications in detail; here we refer to them by name once the context is clear.

The overall thesis is simple: the reserve question is not answered by one object, and the assets that already answer parts of it are not superseded by this one. Alongside them, a dense digital civilization needs a **triad of verifiable necessities**: Privacy, Proofs, and Compute. Treat this triad as services that may support an associated cypherpunk base asset:

- **Privacy**: the right to hold and move value without chokepoints (to preserve agency).
- **Proofs**: cryptographic attestations of origin, integrity, identity, or computation (to crystallize truth).
- **Compute**: useful work (ZK proving, matrix multiplication, verified inference) that applications demand and that blockchains can verify cheaply (to power intelligence).

These are not slogans. They are the cypherpunk stack distilled into services: what cannot be forged, what no one has to bless, and what everyone can check. They are *scarce capacities* the world may continually buy because life in a dense, digital civilization requires them. Each has its own economy. Whether one monetary asset should sit beneath all three is a hypothesis tested below, not a conclusion smuggled into the definition.

Institutionally, this is not just a design for another chain. It is a **research and engineering agenda**: a Bell Labs for proof-of-useful-work and the zk economy. Where the original Bell Labs turned Shannon’s information theory into cables, switches, and semiconductors, the mandate here is to *turn privacy primitives, zero-knowledge, and verifiable compute into everyday infrastructure*: receipts, rails, and verifiable machines that other people build on without thinking about it. This document is the charter for that lab.

Each primitive is designed to be credibly scarce, permissionless, censorship-resistant, and cheap to verify. Gold condenses geology; Bitcoin condenses thermodynamics; this triad condenses the utilities of the information era into assets: things that cannot be faked and do not ask permission, cheap for anyone to verify yet costly to produce or deny. Three different

hard facts, answering three different questions, none of them a later edition of another (Section 31.1.1).

The claim under test. Under sustained administrative and financial repression, a bearer asset may earn monetary premium if holding it preserves private settlement, portable proof, and access to verified compute after ordinary substitutes weaken. The full service path must remain stress-deliverable; demand must not bypass the asset; a persistent, self-custodied, loss-bearing holder constituency must absorb residual financial risk; and project credit must remain separate. Whether any network satisfies those conditions—or whether the triad should share one asset—is an empirical question, and this thesis is written so that the answer can be no.

The thesis will test this through a conditional chain: utility demand creates standardized work; standardized work produces receipts; receipts represent stress-deliverable service; those receipts enable markets; markets produce fee flows; fee flows support a scarce asset *only under non-bypassable accrual rules*; a persistent, self-custodied, loss-bearing, regime-responsive holder constituency anchors that asset; and only then, if the asset remains liquid, neutral, verifiable, legally holdable, and agency-preserving, may it earn a store-of-value premium.

1.2 Premises

The argument below is the steel-man form used on the project site and in outreach. Each premise is developed in the body; together they bound what the thesis does and does not claim.

- P1. Soft guarantees are weakening.** Modern money, media, and infrastructure increasingly depend on compliance systems, custodians, platforms, hardware vendors, and institutions whose guarantees are politically and technically fragile (Chapter 2, Chapter 3).
- P2. The digital economy has three unavoidable needs.** A dense AI civilization needs private settlement, portable attestations, and verified compute—Privacy, Proofs, and Compute as scarce capacities, not slogans (Section 1.1, Part II).
- P3. These needs can become verifiable commodities.** If workloads are standardized, outputs become receipts, and verification stays much cheaper than production, proofs and verified compute can trade as public facts rather than platform promises (Chapter 20, Appendix A).
- P4. Gross capacity is not deliverable service.** Installed power and hardware are inputs, not usable output. Monetary-grade capacity is the scenario-specific surviving flow through energy, conversion, cooling, hardware, network, workload, proof, verification, settlement, and user access (Section 7.2, Section 15.9).
- P5. A store of value requires more than utility.** Indispensable services do not automatically

become money; demand must accrue to a scarce asset rather than leaking to operators, hyperscalers, custodians, or fiat rails (Chapter 4, Section 31.1.4).

- P6. Value capture requires enforceable monetary design.** The asset must be required for core fees, partially burned or retired, staked as collateral, issued under discipline, and protected from bypass channels (Value Capture Lemma, Section 28.6).
- P7. Gross native demand is not a monetary anchor.** Just-in-time fee acquisition, burns, operator inventory, collateral requirements, and wrappers do not by themselves create a persistent, self-custodied, loss-bearing constituency that adds through stress (Section 24.8.8, Section 24.8.7).
- P8. The system must remain falsifiable.** Verification cost, reachability, settlement safety, decentralization, fee coverage, native-asset capture, delivered capacity, holder quality, and agency must be public telemetry—not marketing claims (Chapter 24, Section 28.6).
- P9. Market price is not proof of monetary adoption.** Even when the protocol loop works, observed price can be set by wrappers, leverage, dealers, and allocation rules without native use (Corollary 11.1, Section 24.8). VerifyFlow instruments the second loop; it does not replace VerifyPrice, VerifyReach, VerifySettle, or the Native Monetary Buyer Map.
- P10. Duration-neutral money is not duration finance.** A repression-resistant store of value cannot be a coupon the state can pin. Civilization still needs someone to warehouse the years between pouring concrete and producing power. Project notes are explicit loss-bearing credit, not monetary backing (Section 3.3, Section 31.5.2, Section 31.5.1).

Conclusion (conditional). A base asset is only a monetary candidate. It may earn premium only if the seven-layer stack delivers usable service under stress, demand cannot bypass it, and a persistent self-custodied constituency bears loss through the relevant regime. Work Credits remain typed capacity or service claims; FCR, FER, and PIDL artifacts remain evidence; project notes remain duration-bearing credit; and LP or staking positions remain derivatives or operating claims. No useful service is promoted to money by terminology.

1.3 Contributions

This thesis makes four main contributions:

1. Threat model and layered architecture.

We build a threat model that assumes intentional repression rather than benevolence: financial repression (YCC, capital controls), censorship and shutdowns, hardware and identity capture, and platform-mediated reality. On top of that we propose a seven-layer “cypher-punk stack”:

- Layer 0: Verifiable Machines & Energy

- Layer 1: Reachability (communications & transport)
- Layer 2: Distribution & Execution (software supply & runtime)
- Layer 3: Identity & Claims (humans and machines without doxxing)
- Layer 4: Truth & Work (proof systems, PoUW, VerifyPrice)
- Layer 5: Value & Settlement (privacy rails, non-custodial flow)
- Layer 6: Governance & Telemetry (keeping neutrality and resilience measurable)

The rest of the thesis walks this stack from silicon and power up through proofs, settlement, and governance.

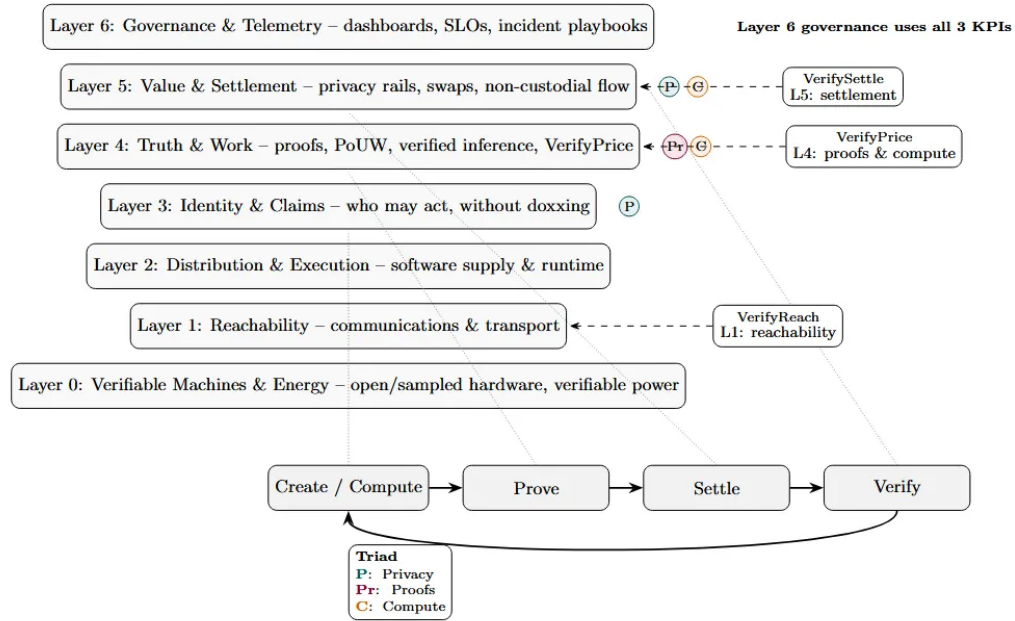


Figure 1: Seven-layer cypherpunk stack with Create/Compute→Prove→Settle→Verify loop, triad overlay (Privacy, Proofs, Compute), and telemetry KPIs (VerifyReach, VerifyPrice, VerifySettle).

Figure 1.1: Architecture Map: The seven-layer cypherpunk stack.

2. Economic formalization: verification asymmetry and VerifyPrice.

We formalize verification asymmetry as the ratio

$$r(W) = \frac{v(W)}{p(W)}$$

between verification and production cost for a workload W , and introduce **VerifyPrice** (a public KPI vector of p50/p95 verify times, costs, and failure rates) as the hinge that turns proofs and verified FLOPs into commodities rather than platform IOUs. This gives a quantitative basis for treating “AI Money” and “ZK Money” as analytical lenses on service

demand, not for treating Work Credits as money. The associated base asset remains only a conditional monetary candidate.

3. **Modular stack: four reference applications and twelve primitives.**

We propose a **Create/Compute Prove Settle Verify** loop and instantiate it with four reference applications (private treasury & payroll, media provenance & authenticity, verified inference, and proof/compute procurement), built on a reusable kit of primitives. These include a Proofs-as-a-Library SDK (PAL), a Privacy Rails Kit (PRK) for non-custodial settlement, Proof Interface Definition Language (PIDL) as a minimal receipt schema, MatMul-PoUW and verified-inference harnesses, canonical workload registries, multi-ZK adapters, SLA escrow/slashing, neutral routers, bridge-safety templates, and a telemetry layer that keeps useful work and neutrality measurable across chains and vendors.

4. **Telemetry, governance, and implementation playbook.**

We extend VerifyPrice into a broader observability regime, including **VerifyReach** (reachability under censorship) and **VerifySettle** (settlement success and refund safety), and propose a “no dashboards, no trust” governance posture where protocol changes and incident responses are driven by SLOs and public receipts rather than foundation fiat. We complement this with an adoption curve, an operator/investor checklist for SoV evaluation, and implementation sketches (Layer-0 hardware, privacy corridors, proof factories, developer SDKs) that make the stack actionable for builders and allocators over the next 12–36 months.

1.4 End-to-End Vignette: The Loop in Action

Before diving into theory, here is a concrete story that walks the **Create/Compute Prove Settle Verify** loop. This vignette shows how the stack works as a felt reality, not just a framework.

Scenario: A Company Runs Private, Auditable Payroll

TechCo is a software company with employees and vendors across several jurisdictions. Its CFO needs to run payroll while preserving salary confidentiality, avoiding unnecessary exposure of the transaction graph, and giving auditors and tax authorities scoped evidence that payments were authorized, complete, and properly reported.

Step 1: Create (Intent) The CFO opens TechCo’s treasury dashboard (built on PRK, the Privacy Rails Kit). She creates a payroll batch:

- 47 employees
- Total: 142,000 USD-equivalent in Work Credits
- Policy: “Salaries are private; aggregate spend is auditable; individual amounts disclosed only with employee consent.”

The dashboard compiles this into a **claim**: “Pay these 47 recipients the specified amounts, under this policy, by end-of-day Friday.”

Step 2: Prove (Compliance + Privacy) The claim is sent to PAL (Proofs-as-a-Library). PAL compiles it into two proofs:

1. **Compliance proof:** A ZK proof that the batch satisfies TechCo’s internal policy (aggregate under budget, recipients are on the approved list, no single payment exceeds threshold). This proof reveals *nothing* about individual amounts or identities only that the rules were followed.
2. **Integrity proof:** A hash commitment binding the full payroll data. This hash is stored on-chain; the actual data stays encrypted in TechCo’s vault.

Both proofs are packaged into a **PIDL receipt**: claim hash, proof hashes, workload ID, SLA tier, timestamps, and a prover signature.

VerifyPrice checkpoint: The compliance proof took 2.3 seconds and cost \$0.004 to generate. Independent verification takes 0.8 seconds on a laptop. This is well within the SLO ($p_{95,t} \leq 5s$, $p_{95,c} \leq \$0.01$).

Step 3: Settle (Private Execution) With proofs in hand, PRK executes the payroll:

- For employees in jurisdictions with limited banking access: atomic swap via **BTC↔XMR corridor**. TechCo’s BTC is swapped for XMR, which is sent to employee wallets. The corridor is non-custodial; if anything fails, funds return to TechCo (refund-safe).
- For employees elsewhere: direct settlement over a **shielded pool** (e.g., Zcash). Each payment is encrypted; only the recipient and TechCo (via viewing keys) can see the details.

VerifySettle checkpoint:

- Swap success rate: 98% (2 retries due to network latency, both succeeded; counted in the 98%, so 49 of 50 corridor payments landed within two attempts).
- Refund safety: 100% (all potential failures would have returned funds).
- Time to finality: median 4 minutes, p95 11 minutes.
- Anonymity set: 12,000+ active notes in the shielded pool.
- These figures are illustrative of the checkpoint’s shape, not observed telemetry; the published VerifySettle schema is the same one the Settlement & Privacy Board reports ([Section 24.9](#)).

Step 4: Verify (Audit Trail) After settlement, the following are publicly verifiable:

- **On-chain:** The PIDL receipt exists, the proofs verify, the aggregate amount was transferred at the specified time.
- **By TechCo’s auditors:** Using viewing keys, auditors can see the full breakdown (who got paid how much) and confirm it matches the compliance proof.

- **By employees:** Each employee can verify their own payment arrived, using their private key.
- **By anyone:** The compliance proof demonstrates policy was followed, without revealing any private data.

What the adversary sees:

- An external observer monitoring Country A sees: “TechCo moved some BTC into an XMR corridor.” They cannot see amounts, recipients, or purposes. TechCo can provide auditors and tax authorities scoped viewing keys and PIDL receipts demonstrating that payments were authorized and properly reported.
- A competitor monitoring the blockchain sees: “Some shielded transactions occurred.” They learn nothing about TechCo’s payroll structure or employee compensation.
- A hacker who compromises a single employee’s device sees: that employee’s payment. They cannot reconstruct the full payroll or identify other employees.

What TechCo gains:

- **Privacy:** Payroll data is not exposed to competitors, data brokers, or unrelated intermediaries by default. Auditors and tax authorities receive scoped evidence via viewing keys.
- **Compliance:** Auditors get cryptographic proof that policy was followed, without needing to trust TechCo’s word.
- **Resilience:** Even if banks freeze accounts or exchanges delist, the privacy corridor remains operational.
- **Receipts:** Every step produced a verifiable artifact (PIDL receipts, proofs) that can be archived, audited, or used in disputes.

Stage	What Happens	Output
Create	CFO defines intent + policy	Claim
Prove	PAL compiles compliance + integrity proofs	PIDL receipt
Settle	PRK executes via privacy corridors	Value transferred
Verify	Anyone can check proofs; auditors use viewing keys	Audit trail

Table 1.1: The Create/Compute Prove Settle Verify loop in the payroll scenario.

The loop in summary: This is **lawful privacy**: default-private, optional-disclosure, with receipts that anyone can verify. The triad (Privacy for settlement, Proofs for compliance, Compute for proof generation) works together to make the flow possible.

Variation: Media Provenance

The same loop applies to a journalist publishing a video:

1. **Create:** Camera with secure element captures footage, emitting a signed provenance attestation (device ID, timestamp, geolocation hash).
2. **Prove:** Edit suite issues proofs of each transformation (crop, color grade, caption). Each edit is a new PIDL receipt referencing the parent.
3. **Settle:** (Optional) If the video is monetized, payment flows over privacy rails to the creator, with receipts linking payment to provenance.
4. **Verify:** Any viewer or platform can verify the chain: “This video originated from camera X at time T, was edited as follows, and has not been tampered with since.” Deepfake detectors can check against registered provenances.

Platforms that strip metadata cannot strip the on-chain PIDL receipt. The proof persists even if the platform delists the content.

These vignettes are not hypothetical futures; they are the concrete scenarios the rest of this thesis is engineered to support. The stack exists to make these flows **cheap, verifiable, and non-custodial** under adversarial conditions.

Chapter 2

The Failure of Soft Guarantees

The twentieth century ran on soft guarantees: the promise that money would preserve purchasing power; that televised images would correspond to facts; that institutions would be slow, sober, and legible. The twenty-first runs on different physics. Monetary systems must service accumulated promises faster than real productivity can grow. Media systems can synthesize plausible realities faster than editors can adjudicate them. Intelligence, once gated by human attention, now scales with silicon. In that environment, guarantees that depended on slow institutions and centralized editors begin to fray.

One soft guarantee is failing without fraud and without default, which is why it belongs in this list rather than in a chapter about mismanagement: the behavioral promise that discipline converts into advancement — gain skills, borrow for education, save regularly, buy the index, purchase a home, and compounding delivers security. Every contract in that recipe can be honored exactly as written while the mapping fails, because the incumbent’s balance-sheet wealth is the entrant’s acquisition price, and the financing regime that produced the earlier returns has already closed (“The Vibecession Was Real”; Sources). A trust failure that survives audit is the hardest kind to repair, and it feeds the participation line directly (Section 3.6).

Since the end of Bretton Woods, *compliance*, not convertibility, has been the hidden backing of money. When compliance becomes weaponized and asset freezing becomes policy, savings need privacy the way ships need hulls.

*When the **symbols of trust** are cheaper to counterfeit than to maintain, the market seeks a new substrate.*

This is not moral decline; it is entropic arithmetic. When verification is costly and authority is centralized, cheaters arbitrage the gap. When verification becomes cheap and public at scale, the equilibrium flips. Our civilization is now crossing that threshold.

“Money is memory with consequences.”

It records who did work and who is owed work, then carries that record across time, space, and politics. Gold solved this with geology; fiat with law, taxation, and war; Bitcoin with thermodynamics and code. Each regime answered the same question (what counts as real work, and how do we know?) with a different ontology.

2.0.1 A Prior Transition: What the Press Actually Changed

Each of those regimes also rested on a medium, and each medium restructured the trust that ran on it before anyone noticed. This is the pattern the media-transition literature documents—as mechanism, not measurement—and it is worth one paragraph here because the thesis’s claim about attestation has a predecessor with a five-hundred-year record.

Manuscript culture anchored trust in persons: the scribe was known, the chain of copying was personal, and verification meant asking a community that remembered. The printing press collapsed the cost of *replication* and, in doing so, moved the anchor from person to artifact: the imprint, the uniform edition, the publisher’s mark. What a reader could check changed—from reputation to internal consistency and corroboration across copies—and the standardization that made every letter identical to every other letter of its kind later gave industry its template for interchangeable parts (Eisenstein 1979; Sources). The broadcast era then re-centralized attestation: one attester, many receivers, and verification deferred back to the institution.

The networked era is the first with no native anchor at all. Anyone can publish; nothing arrives pre-vouched. The incumbent answer is the platform label—attestation by intermediary, and therefore by whoever pressures the intermediary. The answer this thesis tests is attestation by proof: the anchor moves from person (manuscript), to artifact (print), to computation itself. The axis is not new; what is new is that the checking is native to the medium rather than borrowed from the institutions that grew up around the previous one.

Two cautions keep this framing honest. The press’s consequences arrived over centuries that nobody planned, which is an argument for instrumenting rather than forecasting, not an argument that any particular consequence follows. And the literature itself is synthesis rather than experiment; it is used here for pattern, not for quantities.

Memory alone is not money; the record must attach to a scarce object or claim whose holders capture economic value. The triad thesis must therefore show not just that these capacities produce records, but that a specific asset captures the economics of those records.

Our era answers: **work is what machines can do and humans can verify cheaply**. The new memory must therefore be backed by capacities essential to digital life and secured by proofs anyone can check. The unit must survive not because a state decrees it, but because the network will keep buying its utility through booms, busts, and repression cycles.

“A store of value is not a museum; it is a power plant that stays online.”

The same forces that erode the old legitimacy (AI, ubiquitous networks, programmable markets) also supply the antidote. Cryptography can wrap data, identity, and computation in attestations that travel; privacy tech can re-sacralize the personal and the intimate; verifiable compute can make expensive work legible to cheap verification. The system that dissolves

trust also furnishes the reagents to precipitate it again, this time as math, not authority. Verification, not violence, becomes the final arbiter of bounded claims about origin, computation, and settlement.

In the rest of this thesis we make that sketch precise. We start from a threat model that assumes repression, not benevolence: debt stocks that make financial repression arithmetically attractive; a communications stack that can be filtered or shut off; hardware and identity systems that can be integrated into surveillance infrastructure. Against that backdrop we propose a layered architecture, from verifiable machines and energy up through communications, distribution, identity, truth, settlement, and telemetry, and show how three capacities (Privacy, Proofs, and Compute) can be engineered across those layers as services capable of supporting a conditional base-asset monetary candidate. The task is not to find a new story to believe in, but to build and measure the full chain.

During the transition, old-world hard assets—gold, silver, and other geologically scarce commodities—remain rational hedges against the very soft-guarantee failures we describe. They require no counterparty, no network uptime, no software update. The question is not *whether* to hold them, but *when* the new stack becomes robust enough to migrate. Until VerifyPrice dashboards are live, privacy corridors are refund-safe, and receipt volume is non-trivial, the bridge assets remain load-bearing. This thesis is not a call to abandon them prematurely; it is a blueprint for what comes next.

Chapter 3

The World Forces New Monetary Primitives

The forces doing the forcing run on two channels at once, and they are the same channel. In truth, the incumbent order is broadcast: one attester—state, rating agency, platform, exchange—asserts, and N receivers accept. In value, the same shape: central-bank credit is one issuer to N accounts, with the sovereign acting as the typesetter of uniform value, every unit identical because one authority struck it. Proofs and bearer settlement are both $N \leftrightarrow N$: anyone checks anyone's claim; anyone settles with anyone directly. The thesis is one observation wearing two dresses—truth and value move on the same topology, and both are moving from broadcast to network. That is a way of reading what follows, not an additional premise: the chapters that follow treat macro pressure (Section 3.1) and epistemic pressure (Chapter 2) as one phenomenon observed from two sides, and each must still carry its own evidence.

3.1 Macro Playbook: Debt Repression Flight to Neutrality

Total global debt (private + public) according to the IMF's latest Global Debt Database now sits at a roughly $\sim 2.4\times$ multiple of world GDP, with some economies (such as Japan) exceeding 400% of GDP when public and private debt are combined. Standard scenarios from multilateral institutions point to global public debt approaching $\sim 100\%$ of GDP by decade-end, after already exceeding \$100 trillion in the mid-2020s.

Historically, such overhangs are closed via **financial repression**: inflation plus regulation that imposes negative real returns, including overt capital controls and captive-balance-sheet rules, pushing savers toward neutral, censorship-resistant rails. “Voting with your feet” is harder when the exits are policed at gunpoint.

The Bondholder Kill Box (How Repression Works)

Definition. Financial repression is a stealth tax on savings: cap nominal rates, let inflation run, and force captive balance sheets to hold the paper. The result is systematically negative real returns for depositors and bondholders; an engineered wealth transfer

to the sovereign. In the post-WWII advanced economies, real rates were negative in roughly *half the years from 1945–1980*, and the “liquidation tax” (interest-expense savings on the public debt) averaged on the order of 1–5% of GDP per year.

How the cap is enforced. Overt or implicit yield-curve control (YCC), administered rate ceilings, capital/liquidity rules that force government paper into banks, insurers, pensions, and collateral boxes, plus capital controls to slow leakage. In the U.S. WWII episode, the Fed capped T-bills at 3/8% and long bonds at 2½% until the 1951 Accord; when the cap ended, long-duration holders took capital losses as yields normalized.

Why now. With global debt at several times world GDP, public debt ratios trending toward ~100% of GDP by decade-end, and interest costs compounding, repression is arithmetically attractive to policymakers versus explicit default or politically costly austerity.

Two equations every saver should know:

1. Repression wedge (real carry):

$$r_{\text{real}} \approx \frac{1 + i_{\text{capped}}}{1 + \pi} - 1$$

If $i_{\text{capped}} = 2.5\%$ and inflation $\pi = 6\%$, the real return is

$$\frac{1.025}{1.06} - 1 \approx -3.3\% \text{ per year.}$$

2. Duration burn (when the peg breaks): For a bond with duration D , a yield move Δy produces a price move

$$\Delta P \approx -D \cdot \Delta y.$$

A 10-year with $D \approx 8$ takes about -24% on a $+300$ bps repricing.

A worked example: three years at -3.3% real/yr is $\approx -9.6\%$ real loss compounded; then a $+300$ bps jump deals another $\approx -24\%$ price loss, roughly -31% cumulative in purchasing-power terms. The design is deliberate: bleed bondholders while buying time for the sovereign balance sheet.

3.1.1 The Treasury Market as Control Panel

Modern repression does not need to announce itself as yield-curve control. It can be routed through balance-sheet architecture. A sovereign that cannot fund itself at market-clearing real yields can manufacture demand for its paper by changing the rules under which banks, insurers, pensions, money-market funds, stablecoin issuers, and custodians hold collateral.

The playbook is subtler than confiscation: lower front-end rates, pressure or tolerate higher long-end yields, steepen the curve, loosen bank leverage treatment for sovereign paper, and let regulated balance sheets absorb debt that private investors would not hold at prevailing real yields. The central bank can claim discipline because its own balance sheet is shrinking, while the banking system performs the functional equivalent of QE. In practice this shows up as redefined high-quality-liquid-asset treatment, eased leverage-ratio exemptions for Treasuries, incentives for banks to intermediate sovereign debt, privileged government collateral in payment and clearing systems, steered stablecoin reserve composition, and regulated custody as the default gateway for institutional capital.

Balance-Sheet Repression

Financial repression implemented through collateral rules, capital treatment, stablecoin reserve requirements, custody mandates, and institutional balance-sheet incentives, rather than through overt yield-curve control or capital controls.

For a post-fiat monetary stack, this matters because the adversary is not only censorship. It is **substitution**. Regulated wrappers, approved custody, stablecoin reserve mandates, and institutional balance-sheet rules can direct capital into price exposure while starving protocol-native privacy, proof, settlement, fee-burn, and collateral loops. The stack must therefore monitor not only whether users *can* access the asset, but whether they are using the monetary rail or merely holding a compliant shadow claim. Section 11.13 formalizes this substitution risk as the **Wrapper Dominance Ratio**.

3.1.2 Who Actually Buys the Paper

A control panel is only as useful as the hands available to absorb what it issues. The composition of those hands has changed, and the change matters more than the level of any particular yield.

Foreign official reserve managers—the patient, price-insensitive creditors of the previous era—have not been dependable net buyers of U.S. government debt for some years. A growing share of marginal demand now comes from leveraged relative-value funds that hold cash Treasuries against short futures positions and finance the difference in the repo market. Federal Reserve staff research estimates that Treasury holdings of Cayman-domiciled hedge funds rose by approximately \$1 trillion after 2022, reaching roughly \$1.85 trillion by the end of 2024, financed through bilateral and centrally cleared sponsored repo.

Two qualifications are load-bearing, and the source states both. Not all of that exposure is the cash–futures basis trade; hedge funds hold Treasuries for several repo-supported arbitrage reasons. And an estimate assembled from regulatory filings is not a market census. What survives both qualifications is structural rather than quantitative:

A meaningful share of marginal demand for the world's reference safe asset is now supplied by leveraged, repo-dependent, balance-sheet-constrained intermediaries rather than by patient reserve managers.

This changes the failure mode rather than the debt level. Patient creditors withdraw slowly, for reasons of policy, and announce themselves. Leveraged intermediaries withdraw abruptly, for reasons of margin, and do not. The distinction yields a definition the rest of this chapter needs:

Sovereign Credit Safety vs. Collateral Stability

Credit safety is the probability that a sovereign obligation is repaid in nominal terms. **Collateral stability** is the reliability of that obligation's price over the horizon at which it is actually used—as margin, as liquidity, and as the reference rate for everything else. An instrument can be entirely money-good and simultaneously an unreliable short-horizon hedge, because the institutions financing it are leverage-sensitive even when the issuer is not.

The two properties are routinely conflated, and the conflation is load-bearing for a great deal of portfolio construction. A saver who holds sovereign paper for credit safety and receives collateral instability has not been defrauded; they have been sold one property and have priced the other.

A measurement note. The same research finds these positions are largely invisible in the Treasury International Capital data, which understate Cayman holdings by roughly \$1.4 trillion at end-2024 because repo collateral transfers break the reported chain of ownership. Because the TIC data feed the Financial Accounts of the United States, and because the household sector's Treasury holdings are computed there as a residual, the error propagates: published household holdings are overstated, and the measured personal saving rate was overstated by roughly 2.1 percentage points during 2023–24.

This is worth dwelling on, because it is the exact pathology this thesis asks protocols to avoid. A structurally important flow grew to trillion-dollar scale inside the most heavily instrumented financial system in the world, and the official accounts went on publishing confident figures about household balance sheets the entire time. [Section 24.7](#) argues that telemetry must be built so unmeasured dependencies surface rather than disappear into a residual. Here the residual was called “households.”

Automatic Demand Is Not an Anchor

The leveraged residual is only half of the buyer-quality problem. The other half is that a large, persistent bid can still fail to warehouse interest-rate risk.

Green's later work on the long end of the Treasury market distinguishes two kinds of buyer. The **automatic buyer** follows a rule—default enrollment, a target-date glide path, a market-value-weighted bond index—and continues to purchase regardless of whether the offered yield compensates the risk. The **marginal buyer** absorbs whatever duration remains after automatic and other captive allocations, and therefore sets the clearing price. Substantial gross demand can coexist with an unstable market if most of that demand is short-maturity, leveraged, one-time, or insensitive to an increase in risk.

Automatic vs. Marginal Buyer

The automatic buyer supplies demand by formula. The marginal buyer sets the price by absorbing the residual. Quantity of bids and quality of risk absorption are different objects, and they are routinely reported as if they were the same.

The replacement of Japanese life insurers and quantitative-easing central banks—institutions with enduring liabilities, or with a policy mandate to hold duration—by American retirement defaults is therefore not a like-for-like substitution. The old buyer bought duration to match a promise. The new buyer follows an index weight.

The insurer had liabilities. The index has weights.

That sentence is the same transition this thesis already names in verification and in retirement architecture more generally: a rule can execute faithfully while producing the wrong aggregate result. Verifiability guarantees that a rule was followed. It does not guarantee that the rule was wise.

A Fixed Dollar Inflow Is Not a Fixed Risk Bid

Market-value-weighted bond indexes allocate each new contribution according to current market value. When a long bond's price falls, its index weight falls, and a smaller share of every subsequent inflow is directed toward it. Existing holdings are not sold by the rule; incremental stabilizing capacity is. The automatic system therefore absorbs **less interest-rate risk per dollar precisely when long bonds have cheapened and the residual is largest**.

The relevant unit is not dollars of demand but **DV01**: the dollar change in value produced by a one-basis-point change in yield. Automatic inflows can remain large in cash terms while DV01 absorption shrinks. The residual then falls more heavily on the leveraged and discretionary buyers already identified above.

This is a compositional analogue of the Market Realization Plane (Section 11.14). In equities, market-cap weighting directs more dollars toward names that have already risen. In bonds, market-value weighting directs less duration risk toward names that have already fallen. In both cases the governing rule is price-insensitive in the ordinary sense and **pro-**

cyclical in risk terms. Target-date rebalancing, which is often cited as a stabilizer, is itself a relative-performance rule: it fires when stocks and bonds diverge, and it contributes little when they fall together. [Appendix I.1](#) traces that case as a scenario.

None of this requires treating Green’s attribution of any particular yield move as load-bearing. The structural claim is narrower and is the one used here: **the quality, horizon, funding stability, and risk-bearing capacity of demand for long-duration sovereign paper have deteriorated, even where the quantity of automatic buying has not.**

The same distinction splits credit risk from duration risk on the corporate side of the AI complex. Investors may still believe a hyperscaler will repay and still refuse to lend for thirty years at the offered yield. The first crack in the collateral loop ([Section 3.2](#)) need not be an equity collapse. It can be the loss of quasi-sovereign status for long corporate paper—a maturity-specific scarcity of buyers, not a scarcity of nominally safe claims.

3.2 The Collateral Loop

The preceding sections describe a sovereign managing the demand side of its own debt market. That description is incomplete in one important way. It treats asset prices as something the sovereign observes and occasionally influences. Increasingly they are something it depends on.

Collateralized Sovereign Stack

An arrangement in which a state’s effective fiscal capacity depends not only on taxation and monetary authority but also on the market value of the asset complexes it regulates—because those values determine taxable financial income, consumption through the wealth effect, foreign capital inflows, private credit creation, and the collateral against which the private sector borrows.

This Is Not an Eighth Layer

The collateralized sovereign stack is a description of the incumbent order. The seven layers of [Chapter 6](#) are a specification for something to be built. They are different kinds of object and are not to be merged: no layer of the cypherpunk stack produces equity valuations or repo collateral, and no amount of Layer 0 capacity substitutes for a functioning sovereign balance sheet. The duration warehouse of [Section 3.3](#) belongs in the same diagnostic column. The parallel drawn in this section is diagnostic, not architectural. The same caution applies here that [Section 11.14](#) applies to market realization: some things are properly modelled around the stack rather than within it.

3.2.1 The Circuit

Four dependencies, each individually unremarkable, close into a loop.

1. **Asset values support receipts.** Capital-gains realizations, taxes on stock-based compensation, corporate profits, and consumption financed by the wealth effect are all functions of market levels rather than of output alone.
2. **Receipts support fiscal capacity.** A weaker receipt base widens the deficit and increases the quantity of debt that must be placed, independent of any change in policy.
3. **Fiscal capacity and central-bank support sustain Treasury market functioning.** The reference asset stays reference-grade only while it can be financed, cleared, and repo'd without incident.
4. **Treasury yields sustain asset values.** The long yield is the discount rate against which every other claim is priced, and Treasuries are the collateral beneath most private credit.

Step 4 returns to step 1. The system is reflexive, and the reflexivity is not a pathology introduced by bad policy—it is the ordinary structure of a financialized economy with a large sovereign balance sheet.

The state has become partially dependent on collateral values it is nominally free to let the market determine.

3.2.2 Why This Explains Balance-Sheet Repression

Section 3.1.1 describes repression routed through collateral rules, capital treatment, reserve mandates, and custody defaults, and asks why a sovereign would choose that route over the blunter instruments. The loop answers the question.

Every classical adjustment mechanism attacks the loop somewhere. Austerity cuts receipts by cutting activity. Higher policy rates raise the discount rate and depress collateral values directly. Explicit yield-curve control announces the constraint and invites a test of it. Default is unavailable to a sovereign that borrows in its own unit.

Balance-sheet repression is attractive precisely because it is the only lever that manufactures demand for sovereign paper *without requiring collateral values to fall*. It operates on who is obliged to hold the debt rather than on what the debt is worth. That is not a marginal preference; it is close to a dominant strategy, and it is the reason to expect this specific form of repression rather than its historical predecessors.

For a monetary stack the consequence is the one already stated in Section 3.1.1: the adversary is substitution rather than prohibition. The loop tells us why substitution is the instrument of choice.

3.2.3 A Branch in the Playbook

Section 3.1 runs from debt through repression to a flight toward neutral, verifiable rails. Nothing in this section changes that direction. It adds a branch at the transition.

The playbook implicitly assumes the standard crisis sequence, in which risk-off lowers long yields and buys the sovereign time. Section 3.1.2 gives a reason that sequence can invert: when marginal demand is leveraged, a volatility shock converts buyers into forced sellers, and the flight to safety can be followed by a rise in long yields rather than a fall. Conventionally safe and conventionally risky assets then fall together.

The branch matters for savers rather than for forecasters. In the standard sequence, duration is the hedge. In the inverted one, duration is a second position in the same risk, and the only assets that behave differently are those whose value does not derive from someone else's balance sheet. That is an argument about the *structure* of neutral collateral, and it is the same argument Chapter 4 makes on other grounds. Appendix I.1 traces the mechanism in full; consistent with the epistemic policy of this document, it is treated there as a scenario rather than a projection.

3.2.4 Why AI Sits at the Center of the Loop

The loop would be an abstraction if the assets inside it were diversified. They are not. One complex currently touches every step, which is what makes a sectoral repricing a sovereign question rather than an investor's question. The roles below are landing zones within the incumbent macro order—where the consequences of the AI complex show up—and not layers of the architecture this document specifies. Several carry names that resemble layers of the cypherpunk stack because the same physical constraints bind in both places, which is a fact about electricity and silicon rather than a claim about architecture.

Two conclusions follow, and they point in opposite directions.

The first is defensive. A significant repricing of this complex would not be contained to technology equities. It would weaken collateral, receipts, consumption, and foreign inflows simultaneously, and it would do so while the mechanism in Section 3.1.2 was making sovereign paper a less reliable hedge. Appendix I.2 traces that chain.

The second matters more for this thesis. If the complex is fiscally load-bearing, the state acquires a strong interest in preventing its repricing—and the instruments available for that purpose convert private infrastructure into protected public infrastructure. Section 5.2 treats this as an adversary class, because from the standpoint of an open stack it is one. The danger to a neutral, verifiable monetary rail is not that the incumbent AI complex fails. It is that the incumbent AI complex is rescued, and the terms of the rescue determine who is permitted to compute.

Role	Where it lands
Electricity, cooling, and interconnection demand	Power and grid capacity
Data centers, chips, transformers, construction	Industrial output and capital expenditure
Models, inference, agents	Compute and communications
Public and private valuations, vendor financing, leases	Collateral and private credit
Capital-gains and stock-compensation receipts	The sovereign balance sheet
Index concentration, passive flows, debt issuance	Market plumbing
Synthetic media and the provenance problem	Verification
Eligibility, surveillance, administrative decision-making	Governance and participation
Long-lived plants, grids, fabs, and data centers	Duration warehouse (Section 3.3)

Table 3.1: One complex, acting simultaneously at every step of the loop. Concentration of this kind is what converts a sectoral repricing into a fiscal event. The last row is the one this section had not named: the same physical assets that load the collateral loop must be financed across decades, and the buyer architecture of Section 3.1.2 is becoming less able to hold that interval.

Sovereign maturity transformation. If private buyers will not absorb long duration at politically acceptable yields, the least overtly coercive immediate option is to issue bills and accept rollover risk. That does not extinguish duration. It transfers it from investors onto the sovereign’s refinancing calendar. A state that finances long-lived infrastructure, defense, or industrial investment with short-term public liabilities has performed, at civilizational scale, the maturity mismatch that periodically destabilizes banks. Gromen’s fiscal-dominance argument describes where that path ends. Green’s buyer-quality argument describes why it begins: the private duration warehouse is inadequate. Neither claim is load-bearing for the red lines. Both belong in the diagnosis of the incumbent order.

3.3 Two Kinds of Duration

Duration-neutrality is a first principle of the *monetary object* (Chapter 4). It is not a complete account of time.

There are two different objects, and they have been sharing one word.

Duration of the Claim vs. Duration of the Project

Duration of the claim is the interest-rate exposure of a promised cash flow—a coupon the state can pin, a bond whose real return repression can drive negative. A repression-

resistant store of value must not be this instrument.

Duration of the project is the years between committing present resources and receiving the output of a plant, grid, fab, reactor, or data center. Someone must warehouse that interval. A rule that allocates by index weight is not that someone.

The current reserve-asset system uses one instrument, the long Treasury, for both functions: politically usable collateral *and* the conversion of present savings into long-lived public and private capacity. Green's long-end analysis shows the second function weakening while the first is still treated as if it were intact. Gold may be superior as politically neutral collateral because it is not another sovereign's liability and has no refinancing requirement. It does not, by itself, finance a reactor. Bitcoin is in the same position on a digital rail. Proofs can make construction progress, energy receipts, covenants, and restructuring states checkable. They cannot delete time risk, inflation risk, construction risk, or political risk.

Neutral settlement is not a substitute for capital formation. A post-fiat system that cannot finance duration is a savings technology, not a complete political economy.

The design instruction is therefore an unbundling, not an expansion of the token:

- **Reserve collateral** should remain duration-neutral, politically hard to pin, and bearer-capable.
- **Duration finance** should remain labeled credit: underwriting, covenants, maturity transformation, loss-bearing equity, and institutions that can survive marks.
- **Proofs** audit the credit. They do not become the bond.

A post-fiat order that cannot tell collateral from credit will recreate opaque leverage. A post-fiat order that refuses credit altogether will not build Layer 0. The completeness objection is stated and answered in [Section 31.5.1](#). The closed-stack comparison in [Chapter 30](#) names the existing answers: commanded balance sheets in one system, retirement defaults and leveraged intermediaries in the other. Neither is a cypherpunk layer.

This Is Not a Tenth SoV Requirement, and It Is Not a Layer 7

Time is a function of the political economy. Making it a property of the monetary object would undo duration-neutrality. The collateralized sovereign stack remains a diagnosis; Layers 0–6 remain a specification. The duration warehouse sits in the diagnosis, beside the collateral loop, not inside the stack.




Figure 3.1: U.S. Treasury bill peg and long-bond cap under WWII-era yield-curve control (stylized). Peg at $3/8\%$ on bills, cap at 2% on long bonds; dashed line marks the 1951 Treasury–Fed Accord. Source: Rose (2021), Federal Reserve Bank of Chicago.




Figure 3.2: Stylized distribution of ex-post real short-term interest rates in advanced economies, 1945–1980. Real rates are negative in roughly half the years, consistent with Reinhart & Sbrancia’s estimates of the financial-repression era.

3.4 The Repression Playbook, Then and Now

The lesson is instructive: as convertibility receded, the compliance and regulatory perimeter became a more important part of fiat's enforceability, especially under capital controls, sanctions, and financial repression. When repression becomes the spread, capital quietly rotates to neutral, bearer-like rails. Some investors and institutions have interpreted reserve seizures and sanctions—notably the freezing of Russia's reserves in early 2022—as reasons to reassess neutral collateral and privacy-preserving settlement.

If “safe” nominal assets become de facto taxes on savings, rational capital routes (not protests) to assets whose issuance cannot be decreed and whose verification is cheap and public.

In that environment, a credible store of value must:

- Avoid being a duration instrument whose real return can be pinned negative by policy; and
- Live on rails that cannot easily be gated by a single jurisdiction.

This is where the triad enters:

- **Privacy** as an anti-seizure hull that makes savings bearer-like again.
- **Proofs** as cheap public verification that replaces platform vouching.
- **Compute** as a non-coupon revenue base whose clearing price floats with nominal budgets rather than being pegged by decree.

Each primitive is designed to remain neutral and verifiable even when traditional financial infrastructure is weaponized.

3.5 Social: The Web's Trust Default Has Flipped

The same technological forces reshaping money are also reshaping trust.

The web used to borrow its epistemic norms from broadcast: seeing was believing, and the job of editors was to gate what got seen. Deepfakes and generative models invert that. AI systems can now synthesize faces, voices, and scenes that are indistinguishable from genuine footage to lay observers. Exposure itself doesn't inoculate people; survey data across multiple countries suggests prior exposure to deepfakes can actually increase susceptibility to misinformation.

In parallel, we get the **liar's dividend**: once everyone knows deepfakes are possible, any inconvenient real video can be dismissed as fake. The result is a crisis of knowing, not just a rise in error rates.

This isn't entirely new (authoritarians have been editing history since Stalin airbrushed Trotsky out of photos), but the cost curve and scale have changed. Deepfake tools now let anyone with a laptop fabricate a leader conceding an election, a riot that never happened, or “footage” of a historical event that subtly rewrites who was present and who was not. The

fear is not just short-term hoaxes; it is revisionist history at machine speed, where archives, livestreams, and “receipts” themselves become contestable. In that world, the substrate of trust shifts from *memory* (“I saw the clip”) to *mechanism* (“I can verify how this artifact came to be”).

At the same time, the feeds themselves are no longer obviously “neutral pipes.” Investigations and hearings around government–platform coordination have documented extensive relationships between agencies, NGOs, and major platforms for content moderation in the name of combating misinformation and foreign influence. Whether one views this as necessary harm reduction or problematic overreach, the empirical point stands: the old story that your feed is simply “what’s popular” no longer survives contact with the evidence. What you see (and don’t see) is the result of political, institutional, and algorithmic priors you don’t control.

Content authentication standards such as C2PA and watermarking prototypes exist on paper, but real-world implementations are voluntary, inconsistent, and often invisible to end-users. Tests across major platforms such as the Washington Post’s 2025 C2PA test show provenance metadata being stripped in transit, and the few disclosures that do survive are buried in UI that most users never touch.

In other words, **the trust default has flipped**. The practical baseline for online media is now “untrusted unless proven,” and the “proven” part cannot safely be left to platform labels or government–platform partnerships. Platform UX is inconsistent; cryptographic provenance and computation proofs are the scalable, cross-platform backstops.

In a SoV context, this matters because monetary systems sit on top of communication systems: if claims about ownership, origin, and behavior are cheap to fake and expensive to audit, then both money and memory become soft targets. Anchoring value in *receipts that anyone can verify*, rather than narratives that someone must vouch for, is the only stable equilibrium.

Recent work constructs PoUW for arbitrary matrix multiplication with verification that is asymptotically cheaper than production ($O(n^2)$ verify vs. $O(n^3)$ produce, via Freivalds-type randomized checking; $1 + o(1)$ denotes overhead that vanishes *relatively* as problem size grows — a constant-factor gap remains at finite n , so “asymptotically” is doing real work here and the finite- n constants are what VerifyPrice measures). Matrix multiplication is the operation that bottlenecks modern AI. This is the asymmetry PoW always needed: hard to produce, cheap to check.

When verification is very cheap relative to production, **honesty becomes a market equilibrium** and commodities emerge (standardized proofs, verified FLOPs) that can be priced, saved, and eventually used as collateral and stores of value.

3.6 Political: The Participation Line Has Broken

Financial repression becomes politically durable when enough citizens lose the capacity to participate. The relevant threshold is not only poverty, but **participation**: the level of redundancy at which a household can fail, retry, move, learn, transact, save, form families, and refuse coercive terms.

Participation Line

The household or organizational threshold below which a person or firm lacks the redundancy to act freely across time—to fail, retry, transact, move, learn, refuse coercive terms, or survive shocks.

Below that line, legal rights remain on paper while agency disappears in practice. Citizens become managed objects: eligibility files, credit scores, benefit cliffs, tax audits, healthcare forms, student-debt balances, real-name accounts, and platform identities. The state does not need to ban exit if every practical route requires permission.

The line is also approached gradually, and the gradient is where the political danger concentrates. A household can remain formally inside the system — banked, employed, housed, consuming — while becoming progressively less able to accumulate: employed but not in a viable career, housed but unable to save, consuming through debt, technically solvent but one shock from default. Aggregate statistics register none of this until households cross an actual delinquency or unemployment threshold, because the aggregates are dollar-weighted and the deterioration is person-weighted (“The Vibecession Was Real”; Sources). Call the approach path the **exhaustion gradient**: a household above the participation line today can be on a trajectory that reaches it, and the aggregates will look resilient the entire way down. This is why participation telemetry must report buffers and trajectory, not just stocks.

In that world, lawful privacy is not romantic secrecy. It is the technical form of the fresh start. Selective disclosure, private settlement, receipt-based reputation, and identity without doxxing are the digital equivalents of bankruptcy, homesteading, and due process: they preserve the option to act without turning the person into a permanent dossier.

Agency-Preserving Infrastructure

Infrastructure that expands a user’s capacity to act without converting the user into a dossier, dependency object, or platform account.

This is why Privacy, Proofs, and Compute are not merely institutional tools. They are agency-preserving infrastructure. The stack must **equip** citizens and firms to act; if it merely manages them more efficiently, it has reproduced the failure it was built to escape. [Section 33.4](#) returns to this as the thesis’s closing civic frame, and [Item 9](#) makes it a formal SoV requirement.

Chapter 4

First Principles: What a SoV Must Survive

A credible store of value must endure time, space, and politics. In practice, that means:

1. **Credible scarcity:** issuance cannot be tweaked at will.
2. **Cheap, public verification:** authenticity is verifiable by anyone, not decided by a platform.
3. **Censorship-resistance & portability:** no chokepoints; global movement by default.
4. **Neutrality & permissionlessness:** open access; rules apply equally.
5. **Native demand:** the asset does something indispensable, beyond serving as a symbol.
6. **Lawful privacy by design:** default privacy with optional disclosure (viewing keys, auditable receipts) so regulated actors can comply without re-introducing custodians or surveillance chokepoints.
7. **Duration-neutrality:** no fixed nominal cash flows to peg. A repression-resistant SoV cannot be a duration instrument whose real return can be driven persistently negative or whose price can be easily mass-managed by policy. Value should come from scarce capacity purchased every cycle, not from a promised coupon. This requirement governs the *claim*. It does not abolish the *project*. Plants, grids, and data centers still take decades; someone must warehouse that interval. Duration-neutrality of money is not a substitute for duration finance of civilization, and the two must not be the same instrument ([Section 3.3](#), [Section 31.5.1](#)).
8. **Asset-level value capture (non-bypassability):** the monetary object must be structurally required for fees, staking, collateral, settlement, governance, or priority access. If users can consume the triad through fiat, stablecoins, direct cloud contracts, or custodial APIs without touching the asset, the store-of-value thesis fails.
9. **Agency preservation:** the system should equip users to act privately, verifiably, and non-custodially without becoming a surveillance or dependency layer. Privacy, identity, reputation, and compliance features must be designed as tools in the user's hands: selective disclosure rather than global inspection, receipt-based reputation rather than biographies, and non-custodial settlement rather than managed accounts. A system that makes users legible but not free has failed the political-economy test even if its cryptography is sound.
Each of these conditions is just the "money as memory" idea made operational: if money

is the record of who did work and who is owed work, then scarcity, verifiability, censorship-resistance, neutrality, native demand, lawful privacy, duration-neutrality, non-bypassability, and agency preservation are the properties that keep that record honest across time, space, and politics—and keep the humans behind it free to act.

Thesis: Privacy, Proofs, and Compute can be engineered to meet all nine, and the world now needs exactly those properties.

We can sketch how the triad maps to SoV requirements:

- **Credible scarcity** the base-asset constitutional schedule, reported separately from DVC-bounded Work Credit issuance.
- **Cheap, public verification** **VerifyPrice(W)** dashboards for canonical workloads W .
- **Censorship-resistance & portability** **VerifyReach** + **VerifySettle** metrics: reachability under censorship and settlement success/refund safety across corridors.
- **Neutrality & permissionlessness** decentralization telemetry: house share, geo/ASN distribution, time-to-first-proof.
- **Native demand** Work Credit utilization, proof/compute fee share of the security budget, and application-level usage.
- **Lawful privacy** corridor compliance receipts, viewing-key usage, anonymity-set health (shielded-pool size, churn, volume).
- **Duration-neutrality** no fixed coupons; revenues from priced workloads and triad usage, with explicit “repression beta” rather than fixed-income-like promises.
- **Asset-level value capture** native-asset fee share, burn rates, collateral lockups, and bypass-channel indicators (Section 11.13).
- **Agency preservation** non-custodial usage share, selective-disclosure ratio, dossier-minimization score, and the percentage of flows requiring no global identity.

In a world that will likely choose *stealth default* (negative real yields and capital controls) over explicit default, “store of value” can’t just be a narrative; it has to clear visible, adversarial stress tests. If the whole point of the stack is to survive yield-curve control, on-/off-ramp throttling, and peg breaks, then we should write down the conditions under which it continues to function and accrue value.

The checklist below turns those claims into operator SLOs that treasuries and allocators can actually underwrite and monitor. Here are several repression stress-tests to keep in mind:

- **YCC shock.** 24–36 months of -300 to -500 bps real yields FeeCoverage (retained fees over security budget, per Section 23.6.5; burns excluded) remains $\geq$ its phase-gate threshold, and the canonical tuple’s t_{95} field remains $\leq 5s$ for core workloads on Laptop-class hardware.
- **On-/off-ramp squeeze.** Non-custodial BTC $\leftrightarrow$ XMR/ZEC routes maintain $\geq 95\%$ success with 100% safe refunds; shielded-pool anonymity sets remain large and growing (no col-

Table 4.1: Store-of-Value Requirements vs. Triad Capabilities

SoV Requirement	Privacy (P)	Proofs (Pr)	Compute (C)	Metrics / SLOs
Credible scarcity	Shielded balances; no selective debasement; predictable issuance schedules	Auditable supply without doxxing; proof capacity tied to hardware and energy	Verified FLOPs limited by physical compute; energy-tied issuance	FERs; facility telemetry; VERIFYPRICE
Cheap, public verification	Privacy proofs verifiable by anyone; private paths for ordinary users	Succinct, cheap-to-check proofs; VERIFYPRICE targets $t_{95} \leq 5s$ on Laptop-class reference hardware, $r(W) \leq 0.3$	Bounded verify cost for heavy workloads; MatMul $O(n^2)$ vs $O(n^3)$	VERIFYPRICE; VerifyReach
Censorship-resistance	Non-custodial atomic swaps and shielded pools; harder for censors to see flows	Receipts portable across chains; proofs of inclusion/exclusion expose filtering	Open-admission prover/miner markets; diverse operators	VerifyReach; VERIFYSETTLE
Neutrality	Privacy by default, not by permission; same privacy for all	Anyone can verify; open circuits and PIDL interfaces; no gatekeepers	Useful-work mining with open hardware paths; no single vendor chokepoint	Regional VERIFYPRICE; dispersion metrics
Native demand	Enterprises and individuals need privacy for operations, payments, savings	Provenance and compliance mandates; proofs demanded by AI, finance, regulators	AI workloads, ZK proving as budget line items	Fee volume; Work Credit utilization
Lawful privacy	Viewing keys and auditable receipts; default-private with narrow exceptions	Proof anchors satisfy audits; policy-aware proofs reveal facts, not full data	Compute SLAs with receipts; policy-tagged workloads	Compliance flows; stress-test results
Duration-neutrality	No fixed coupon; revenues track priced capacity; survives regime change	Proof unit pricing floats with budgets; long-run telemetry	Verified FLOPs clear at market rates, not pegs; anchored in infrastructure	Multi-year FER and SLO stability
Asset-level value capture	Fees, collateral, and settlement denominated in the native asset; non-custodial routes dominate	Proof/verification fees and staking denominated natively; no free-riding on public verifiability	Compute fees, collateral, and priority access require the native asset	Native-fee share; burn rate; collateral lockups; Wrapper Dominance Ratio
Agency preservation	Default private settlement; no protocol backdoors	Selective proof of facts without dossiers	Open admission to verified work	Non-custodial usage share; selective-disclosure ratio; dossier-minimization score; % of flows requiring no global identity

lapse in active notes or volume).

- **Peg break.** If global yields gap +300 bps, triad revenues (proofs/FLOPs, privacy rails usage) track buyer budgets; there are no coupon-like revenue shortfalls that turn the asset into a synthetic bond.

Publish these in dashboards; **no dashboards, no trust.**

Early Falsifiers

The red lines developed in [Chapter 28](#) provide the full falsification regime. In brief, the SoV thesis fails if: verification becomes expensive or gated; refund safety breaks; verification monoculture emerges; telemetry is captured; fee coverage collapses; or value capture fails (users consume triad services but asset demand remains weak because services are paid in fiat/stablecoins or value leaks to operators). See [Section 28.6](#) for precise conditions and response protocols.

Later sections turn these stress tests into explicit SLAs and telemetry: VerifyPrice, reachability, settlement success, and decentralization metrics that operators and allocators can track in the open.

With these nine requirements as design constraints, we can now turn to the primitives that might satisfy them. Before naming those primitives, however, we need to be explicit about who will attack this system and at what layers. The next section states that threat model; the one after lays out the layered architecture that the rest of the thesis will fill in. The entire document can be read as an attempt to engineer Privacy, Proofs, and Compute to satisfy this nine-point contract under adversarial conditions.

We can summarize the whole design posture in one line: our North Star is to pay the machine only for work anyone can verify cheaply, on rails that give humans lawful privacy by default.

Chapter 5

Threat Model

Why a threat model? This thesis argues that Privacy, Proofs, and Compute are distinct services that may support a base-asset monetary candidate only when verification is cheap and public, the full path remains deliverable, and holder agency survives pressure. A threat model makes those conditions falsifiable. It defines what we must protect, who is trying to break it, where the trust boundaries lie, and which metrics decide if the system remains credible under pressure. It sits between first principles and implementation so that every later design choice can be traced to an explicit adversary and invariant.

Every architecture is, implicitly, a bet about who will attack it and how. Up to this point we have argued from first principles: what a store of value must survive, and why Privacy, Proofs, and Compute can be engineered to meet those conditions. The next step is to be explicit about the pressure this stack will live under.

The backdrop is already in view. Macroeconomic arithmetic makes financial repression attractive: debt stocks in the hundreds of percent of GDP, negative real rates by policy, capital controls and captive balance sheets as standard tools rather than emergency measures. At the same time, the web’s trust default has flipped from “trusted unless flagged” to “untrusted unless proven” as deepfakes and coordinated moderation make memory itself a soft target. The combination is simple and brutal: states are incentivized to squeeze balance sheets and communications; platforms are incentivized to mediate reality; hardware and networks are increasingly treated as levers of policy, not neutral infrastructure. Against that field, a “next-generation store of value” is not being asked to survive price volatility; it is being asked to survive a world that optimizes for control.

5.1 Scope and Assets

The scope is the full **Create/Compute Prove Settle Verify** loop and its hardware base layer. At minimum, we need to protect four things:

1. **Integrity of receipts & state.** Proofs, provenance, settlement artifacts, and consensus state must not be silently corrupted.
2. **Confidentiality of flows.** Private settlement with optional disclosure, not privacy that evaporates at the first audit request.

3. **Availability & neutrality.** Open admission for provers and users, anti-capture, and non-seizable settlement paths.
4. **Verifiability economics.** Verification must remain cheap in absolute and relative terms; otherwise “anyone can verify” collapses into “someone we have to trust.”

These map directly to the SoV properties enumerated in [Chapter 4](#). The assets at risk are not just balances on ledgers. They are the whole loop and the matter it runs on.

- At the bottom lie **machines**: chips, RNGs, secure elements, TEEs, capture devices, modems, routers, and the power infrastructure that feeds them.
- Up one level sit the **proofs and receipts** that describe what those machines did: provenance attestations, computation proofs, settlement artifacts, audit trails.
- Above that are the **flows of value and obligation** that ride on those receipts: private pay-rolls, cross-chain settlements, inference contracts, collateral arrangements.
- Threaded through everything are the **communications substrate** that carries claims and proofs, the **software distribution channels** that keep clients coherent, and the **identity and key material** that bind actions to actors without doxxing them.

All of these must remain sufficiently intact that savings, truth, and compute can continue to clear when old guarantees fail.

5.2 Adversary Classes

We care about several classes of adversaries:

State-level repressors. Impose negative real yields, capital controls, and hardware mandates; gate on-/off-ramps; require closed attestation; seek to turn banks, clouds, app stores, and IXPs into enforcement arms.

Platform cartels. Strip or obscure provenance; collude on labels; prefer vendor-mediated “trust” and walled-garden verification.

Strategic patrons. Protect fiscally load-bearing firms through procurement, guarantees, power priority, and regulatory shelter, converting private infrastructure into protected quasi-public infrastructure ([Section 5.2](#)).

Economic attackers. Miner/prover cartels, router capture, MEV/censorship, liquidity games on swap corridors, front-running and soft-forking that tilt rules in their favor.

Hardware/supply-chain adversaries. Trojans, biased RNGs, covert debug paths that turn “proof” into theater; opaque TEEs whose attestation keys terminate in unaccountable HSM farms.

Energy & physical-interdiction adversaries. Curtail, ration, or price-discriminate against verification workloads; deny interconnection; withhold long-lead grid equipment; prioritize state or industrial load over third-party proving ([Section 5.5](#)).

Cryptanalytic / PQ attackers. Proof forging or signature breaks; long-horizon post-quantum risk.

UX / vaporware risk. Spec-drift and unverifiable performance claims that corrode credibility, even if the cryptography is sound.

Compositional / mechanical adversaries. No hostile intent at all: locally rational agents whose incompatible operating rules combine into concentration, procyclicality, and false price signals (Section 5.2).

Of these, the primary adversary in this frame is not a cartoon hacker but the **rational sovereign under stress**. A state that must service promises larger than its productive base will reach, as history shows, for the levers it actually controls: interest-rate caps, yield-curve control, capital controls, regulated custody, and mandatory “secure” hardware and identity schemes. It will be tempted to turn banks, clouds, app stores, hardware vendors, and IXPs into enforcement arms: require real-name registration at the edge; throttle or shut down networks in the name of stability; mandate TEEs whose attestation keys terminate in state-controlled HSMs; define “safety” as surveillance.

In parallel, platforms facing regulatory and reputational risk will centralize moderation and provenance: which media is shown, which credentials count, which proofs are recognized. The danger is not that any one actor becomes an obvious villain, but that their combined incentives re-create a soft but totalizing chokepoint architecture.

Beneath that sovereign–platform axis are **economic adversaries**: miners, provers, routers, and liquidity providers who prefer rent to work. They will collude if they can, capture matching engines and orderflow, quietly prioritize their own routes, soft-fork the rules in their favor, or simply withdraw service when it suits them.

If verification becomes expensive or gated, a priesthood of “trusted verifiers” will emerge, and with it all the familiar pathologies of rent extraction and arbitrary censorship.

And below them, like bedrock or landmines, are **hardware and supply-chain adversaries**: fabs and vendors (sometimes complicit, sometimes merely compromised) who can bias randomness, slip trojans into IP blocks, leave “debug modes” wired to secrets, or ship opaque enclaves that amount to remote-controlled kill switches. If the machine can corrupt inputs, leak witnesses, bias randomness, falsify energy receipts, or centralize useful-work markets, the economic claims built on those machines become unreliable—even if the underlying proof system remains mathematically sound.

Administrative Repression

The most probable control regime is not always a ban. It is **administrative convergence**. Banks, platforms, app stores, identity vendors, cloud providers, telecoms, stablecoin issuers, custodians, and regulators converge around safety, convenience, compliance, and fraud prevention.

Each measure is defensible in isolation; together they turn optional rails into de facto mandatory rails.

Administrative Repression

The conversion of formally optional financial, identity, compute, and settlement rails into practically mandatory rails through custody defaults, compliance rules, app-store control, tax treatment, institutional mandates, benefit systems, and platform terms of service.

The stack must therefore defend against soft capture as much as hard censorship. The key question is not only “Can the protocol still run if banned?” but **“What share of real usage remains non-custodial, private, verifiable, and protocol-native when compliant alternatives are easier?”** This is the same substitution risk introduced in [Section 3.1.1](#) and formalized as the Wrapper Dominance Ratio in [Section 11.13](#).

Compositional and Mechanical Adversaries

Administrative repression still assumes intent, however diffuse. There is a further class of threat that requires no intent at all.

Not every threat to monetary integrity is an attacker. Some emerge from locally rational agents operating through incompatible rules. Product sponsors maximize fee revenue. Investors chase themes or rebalance portfolios. Authorized participants arbitrage discounts to net asset value. Dealers hedge swaps and options. Passive funds follow mandates. Trend systems follow prices. Each actor can behave rationally within its own boundary while the combined machine produces concentration, procyclicality, discontinuity, and false price signals.

Compositional Adversary

A failure mode in which no participant behaves maliciously or even irrationally, but the interaction of their operating rules produces emergent mechanical dependence, correlated behavior, recursive leverage, procyclical amplification, and prices that no longer carry information about the underlying system.

The threat here is not a villain; it is the sum of locally rational walls. A repression-resilient monetary architecture must therefore measure not only hostile actions, but emergent mechanical dependence, correlated operating rules, recursive leverage, and the failure of stabilizing flows. This is the adversary class that the Market Realization Plane ([Section 11.14](#)) exists to make visible, and that VerifyFlow ([Section 24.8](#)) instruments.

With this addition the threat model now spans four distinct kinds of failure: **hostile capture** (state and platform coercion), **institutional co-option** (administrative repression and wrapper

substitution), **endogenous mechanical fragility** (compositional adversaries), and **protective enclosure** (the national-champion pathway below). A design that defends against only the first is defending against the least likely.

Enclosure by Rescue

There is a fourth kind of failure, and it is the one this thesis is least naturally equipped to see, because it arrives disguised as a solution.

Section 3.2 argues that certain private asset complexes have become load-bearing for sovereign fiscal capacity. A state in that position does not merely regulate those firms. It has a direct fiscal interest in their continued valuation, and when that valuation is threatened it will reach for the instruments it controls.

National-Champion Pathway

The conversion of private firms into protected quasi-public infrastructure, undertaken because the state's own fiscal capacity or security position depends on their continued operation. Effected through procurement, loan guarantees, tax treatment, energy and interconnection priority, trade protection, restrictions on foreign substitutes, strategic equity positions, and compliance regimes whose fixed costs favor incumbents.

The distinction from the adversary classes above is worth stating precisely, because the mechanism is easy to file under one of them and it does not belong there.

State-level repressors act *against* a target. Administrative repression (Section 5.2) converges on control through many small defensible measures, none of which is a rescue. Compositional adversaries (Section 5.2) require no intent at all. The national-champion pathway is none of these: it is deliberate, it is protective in intent, and its beneficiaries welcome it. Call it **enclosure by rescue**.

The threat to an open stack is not that the incumbent compute complex fails. It is that the incumbent compute complex is saved, and the terms of the rescue determine who is permitted to compute.

Why this is the most dangerous class for this thesis. Every other adversary here leaves the open stack's value proposition intact or strengthens it. Repression makes neutral rails more attractive. Platform capture makes portable provenance more attractive. Mechanical fragility makes verifiable telemetry more attractive. Enclosure by rescue does the opposite: it produces a competitor that is well-capitalized, technically excellent, energy-privileged, regulatorily sheltered, and genuinely useful. It wins on capability and convenience, which is exactly the ground on which Section 30.9 concedes the open stack cannot win.

What it threatens specifically. Priority access to power and interconnection, which is a Layer 0 input (Chapter 15, Section 5.5). Hardware allocation under supply constraint. Regulatory recognition of some attestations and not others, which decides whose proofs count. And, at the limit, the treatment of independent inference as a compliance problem rather than a normal activity—the enclosure risk of Section 5.6 arriving through industrial policy rather than through platform terms of service.

What to watch. The instruments are enumerated in Appendix I.3, and they are observable in public procurement records, subsidy programs, interconnection queues, and export-control rulemaking well before they are visible in market structure.

5.3 Threats by Loop Stage

We can slice the threat surface by stage in the Create/Compute Prove Settle Verify loop:

Create/Compute.

- *Threats:* poisoned models, mislabeled workloads, biased hardware profiles, hidden accelerators.
- *Mitigation:* canonical workload registries, hardware profiles bound to receipts, open benchmarking, attested hardware provenance.

Prove.

- *Threats:* junk proofs, prover cartels, selective service, “house-only” pricing.
- *Mitigation:* multi-ZK adapters, neutral routers with fairness tests, SLA escrows & slashing; publishing entry-latency, top-N share, and geography/ASN distributions.

Settle.

- *Threats:* corridor censorship, refund failure, deanonymization, bridge compromise.
- *Mitigation:* adaptor-signature swap kits, lawful-privacy corridors (viewing keys + auditable receipts), bridge-safety templates, mandatory refund-safety, public **VerifySettle** metrics.

Verify.

- *Threats:* verification-cost creep; opaque clients; “trust our node” monocultures.
- *Mitigation:* laptop-grade verifiers, reference harnesses, verify predicates embedded in SDKs, public p50/p95 telemetry via the **VerifyPrice** observatory.

Telemetry & Governance.

- *Threats:* silent centralization, KPI gaming, governance capture.

- *Mitigation*: open dashboards for VerifyPrice, VerifyReach, VerifySettle; decentralization metrics (house share, geo/ASN spread, time-to-first-proof); incident reports; hard rules about what must be public before changes ship.

Layer 0.

- *Threats*: “trust the vendor” cliff, mandatory closed TEEs, unmeasured side-channels.
- *Mitigation*: open RTL or microarchitectures where possible, lot sampling and imaging where not, SNARK-wrapped attestations, and explicit side-channel budgets.

5.4 Mapping the Authoritarian Playbook

We can sketch several canonical “playbook moves” and the corresponding counter-moves. The table below maps real-world censorship and surveillance tactics to the stack’s countermeasures:

Beyond these network-level moves, the financial and institutional playbook includes yield-curve control, capital controls, mandatory “secure” hardware mandates, platform-mediated reality attacks, and proof priesthods. For each of these macro attacks, the stack provides specific countermeasures that are detailed throughout the document and formalized as testable stress harnesses.

The high-level pattern:

- **Yield-curve control + captive balance sheets** a duration-neutral base-asset candidate alongside Work Credits priced as typed Privacy/Proofs/Compute service claims rather than monetary reserves.
- **Capital controls + KYC choke points** non-custodial privacy corridors (BTC↔ZEC/XMR) that remain viable without centralized exchanges; VerifySettle metrics that make corridor health public.
- **Mandatory “secure” hardware + identity schemes** Layer-0 verifiable machines with open or sampled designs; hardware-profile receipts; identity schemes at Layer 3 that separate accountability from doxxing.
- **Platform-mediated reality + provenance stripping** PIDL-encoded receipts for media and computation; PAL to compile provenance claims into portable proofs that survive platform stripping.
- **Proof priesthods + closed verifiers** laptop-grade reference verifiers; multi-backend proof systems; VerifyPrice observatories that publish verifier costs and failure rates.

The thesis responds by treating neutrality, privacy, and verifiability as **service-level objectives**, not vibes.

Playbook Move	What It Breaks	Counter
DNS/IP/SNI/QUIC filtering & active probing	Reaching verifiers/bridges; provenance fetch	Default to Tor Snowflake/obfs4, refraction networking; pin ECH; fall back to domain-front-free paths.
Domain fronting curtailed by major clouds	“Collateral freedom” paths disappear	Use Snowflake’s WebRTC and ISP-partner decoy routing; keep multiple CDNs with content-addressed updates.
App-store pressure	Mobile distribution	Parallel update rails (direct, IPFS, USB/QR) + detached sigs; keep APK/IPA sideload guides ready.
Real-name + SIM registration	Pseudonymous ops	VC/anon-cred kit; receipt-based reputation; no SIM/face gates for access.
Government-ordered shutdowns	All of the above	Sat/mesh/sneakernet paths, store-and-forward receipts; pre-provisioned peer lists; publish “blackout drills” results.
Backbone/seabed sabotage	Regional isolation	Multi-landing routing, regional proof markets, diaspora relays; monitor cable incidents.
Bitcoin network metadata capture	Wallet/settlement mapping	BIP-324 everywhere; BTC↔XMR/ZEC adaptor-sig corridors with refund-safety; publish corridor telemetry.
Approved-custody defaults	Self-custody and native usage	VerifySettle/WDR telemetry; native-fee share; non-custodial UX.
Stablecoin-only institutional flows	Base-asset value capture	Required fee medium; auto-acquire-and-burn; native collateral.
App-store wallet restrictions	Distribution and privacy	Sideloads, web clients, signed binaries, hardware wallets.
Tax/reporting complexity	Ordinary-user access	Selective-disclosure receipts; tax export proofs.
“Safe AI” model licensing	Open verified compute	Open workload registries; decentralized proof markets.

Table 5.1: Authoritarian playbook moves and stack countermeasures.

5.5 Energy & Physical Interdiction

The playbook above is financial, informational, and legal. It has a physical companion that is easy to omit because it looks like infrastructure policy rather than monetary policy.

Every threat class enumerated so far assumes the machines are running. The adversary corrupts inputs, gates verification, captures governance, strips provenance, or mandates attestation. In each case there is still a facility drawing power, and the contest is over what that facility is permitted to do. A more elementary move is available: decide what the facility is permitted to *draw*.

A state that finds it awkward to ban verification can simply decline to energize it.

The Instruments

Energy interdiction is unusually attractive to a rational sovereign under stress because it is administratively cheap, legally unremarkable, and rarely legible as repression:

- **Curtailment and rationing.** Interruptible-load classifications applied selectively; verification and proving workloads designated non-essential during scarcity.
- **Interconnection denial and queue position.** Not a prohibition, merely a multi-year wait, applied unevenly.
- **Tariff discrimination.** Punitive rate classes for compute-intensive load, justified on grid-stability or environmental grounds.
- **Long-lead equipment scarcity.** Transformers, switchgear, and turbines allocated by policy rather than by price. Nothing is banned; the queue simply never advances.
- **Load prioritization.** Sovereign or national-champion AI build-outs and industrial demand given precedence over third-party proving for the same electrons.
- **Water and cooling constraints.** Permitting used to reach the same result as an energy denial.

Note that none of these requires hostility toward cryptography, and most survive judicial review comfortably. This is what makes the class distinct from the mandates in [Section 5.2](#): there is no order to comply with, and therefore nothing to resist.

Why This Belongs in the Threat Model

Two reasons, one narrow and one structural.

The narrow reason is that [Section 5.1](#) already names “the power infrastructure that feeds them” among the assets to protect, and then never returns to it. The asset is declared and left undefended.

The structural reason is more serious. The thesis makes verification affordability a constitutional target and treats its breach as fatal. Those targets are correctly described as exogenous

to token price. They are *not* exogenous to the conditions under which power and reference hardware can be obtained, and those conditions are partly set by states. An adversary who can move the physical cost of verification can therefore trigger the thesis's own primary failure condition without touching the mathematics, the governance, or the market.

That is a qualitatively different attack from the rest of this chapter. Every other adversary here must break something we are watching. This one adjusts a parameter we had classified as an environmental constant.

Distinguishing Two Exposures

The precise causal path differs between the two halves of the loop, and conflating them produces overclaims in both directions:

- **Verification** runs on reference hardware at the edge in small amounts. Its exposure is not bulk power but *obtainability of unprivileged reference hardware*, plus edge power and connectivity cost relative to local income.
- **Proving and verified compute** consume bulk power. Their exposure is energy price, firmness, curtailment risk, and competition from industrial and state load.

Bulk electricity prices do not directly break “anyone can verify,” because verification is not a bulk-power activity. Export controls and platform lockdown do. [Section 15.10.1](#) develops this distinction, and [Section 15.8](#) gives it a metric.

Response Posture

The stack's answer is the same one it gives everywhere else: make the exposure measurable, then let it bind against issuance.

- Facility Capacity Receipts ([Section 15.7](#)) already record curtailment exposure, fuel mix, backup duration, and jurisdictional power risk.
- Sovereign optionality ([Section 15.8](#)) aggregates those fields into one exposure that can be published and challenged.
- Risk haircuts ([Section 23.6.4](#)) price fragility into Work Credit issuance rather than discovering it during a disruption.
- A dedicated red line ([Section 28.6](#)) retires the thesis if verification affordability becomes a sovereign policy variable rather than a market outcome.

What the stack cannot do is engineer its way out of physics. If a jurisdiction decides that third-party proving will not be energized, the correct response is dispersion and honest reporting, not a claim that cryptography makes the grid irrelevant.

5.6 Explicit Tensions and Boundary Conditions

Two tensions within this thesis deserve explicit acknowledgment. Ignoring them would make the framework less credible; addressing them honestly strengthens it.

Tension 1: Privacy by Default vs. Coercion at the Social Layer

The concern: Lawful privacy proposes “optional disclosure via viewing keys.” But if disclosure becomes mandatory at chokepoints (exchanges, employers, landlords, visa applications), “optional” becomes fiction. The technical guarantee of privacy could be circumvented by social or legal coercion.

What the system can and cannot guarantee:

Layer	Guarantee	Limitation
Protocol	No backdoors, no master keys, no escrow. Disclosure requires the holder’s key.	Cannot prevent a holder from being coerced to share their key.
Architecture	Non-custodial routes exist. Users can transact without intermediaries.	If all practical routes require KYC, non-custodial option may be theoretical.
Corridor design	Multiple asset paths with redundant liquidity in multiple jurisdictions.	If all corridors are choked simultaneously, exit paths narrow.
Telemetry	VerifySettle makes corridor health public; users can see which routes are viable.	Does not prevent coercion; only makes it visible.

How the stack reduces coercion surface:

1. **Data minimization:** By default, nothing is revealed. Coercion must extract keys, not merely subpoena records that already exist.
2. **Decentralized verification:** No single auditor or platform must see all flows. Selective disclosure can be scoped narrowly.
3. **Multi-jurisdiction design:** Corridors, mirrors, and provers are distributed so no single legal regime can compel universal disclosure.
4. **Exit optionality:** Even under local coercion, users retain technical ability to move assets to less-coerced jurisdictions—if corridors remain open.

Honest acknowledgment: Lawful privacy is a *technical* property, not a social guarantee. It cannot prevent a sufficiently powerful adversary from coercing individuals. What it does is (a) raise the cost of mass surveillance (each disclosure requires individual coercion), (b) preserve optionality for those in less-coerced environments, and (c) make the coercion visible through telemetry rather than hidden in platform logs.

Tension 2: Verified Compute vs. Hyperscaler Dominance

The concern: Compute is dominated by a handful of hyperscalers (AWS, Azure, GCP) and chip vendors (NVIDIA, AMD). If verified compute becomes valuable, won't these incumbents simply dominate issuance and markets, making "decentralized" PoUW a fiction?

What the system offers vs. hyperscalers:

Property	Hyperscaler Compute	Triad Verified Compute
Cost per raw FLOP	Lower (economies of scale)	Higher (proof overhead, smaller operators)
Verifiability	Trust the vendor	Anyone can verify receipts
Permissionlessness	Vendor can ban workloads, customers, regions	Open admission; no TOS-based exclusion
Censorship surface	Single legal entity; can be compelled	Distributed operators; no single chokepoint
Receipts	Vendor-issued invoices	PIDL receipts verifiable by anyone

Why decentralized verified compute wins in specific scenarios:

1. **Repression scenarios:** When a hyperscaler is compelled to ban certain workloads or customers, decentralized alternatives remain available. The value proposition is not "cheaper" but "still available."
2. **Proof-heavy applications:** For workloads where *verifiability* is the product (compliance proofs, provenance, audit trails), the receipt is the value—not raw FLOP cost. Hyperscalers don't currently sell receipts; they sell capacity.
3. **Certain procurement regimes:** Governments, NGOs, and enterprises with sovereignty constraints may prefer compute that doesn't route through foreign hyperscalers or vendor-controlled enclaves.
4. **Composability with privacy:** Hyperscaler inference is logged. Verified inference over privacy rails is not. For sensitive workloads (medical, financial, personal), the privacy premium justifies higher cost.

Honest acknowledgment: Decentralized verified compute will not beat hyperscalers on raw cost. It competes on *verifiability*, *permissionlessness*, and *censorship-resistance*. In a benign environment, hyperscalers win on price. In a repressive or high-stakes environment, verified compute wins on trust properties. The thesis bets that demand for these trust properties is structural and growing.

Telemetry response: If hyperscaler concentration in verified compute markets exceeds thresholds (e.g., >50% of VerifyPrice-tracked capacity from 3 vendors), this is flagged in dashboards. The decentralization metrics in §22–23 make this drift visible, not hidden.

AI Homestead vs. AI Enclosure

The hyperscaler tension above is a cost-and-verifiability question. Underneath it sits a sharper political one. AI is the first new frontier since the closing of the physical frontier, but it can be **homesteaded** or **enclosed**. In the homestead path, models, tools, proofs, and compute markets distribute cognitive leverage outward. Individuals, small firms, local institutions, agents, and open-source networks gain new capacity to act. In the enclosure path, cognition is rented from a handful of walled gardens. Users receive outputs but not sovereignty; access is mediated by terms of service, identity, payment rails, jurisdiction, and platform risk.

The triad is the monetary and technical response to enclosure. Privacy preserves the right to transact and coordinate without becoming fully legible. Proofs preserve the ability to verify claims without trusting platform labels. Verified compute preserves access to machine intelligence with receipts rather than vendor promises.

Homestead Ratio

The share of verified compute, proof generation, and AI-service capacity supplied by open-admission, non-hyperscaler, geographically diverse, independently verifiable operators.

An **Enclosure Risk Flag** triggers when more than a threshold share of VerifyPrice-tracked capacity comes from a small number of hyperscalers, closed TEEs, or single-jurisdiction facilities—the same telemetry response introduced above, read through a civic rather than purely economic lens.

The goal is not simply to make AI outputs auditable. It is to keep the AI frontier from becoming a tollbooth on human agency, the same demand this thesis names **agency preservation** in Section 3.6 and Item 9.

Tension 3: Institutional Holdability vs. Native Monetary Function

The third tension is the one most likely to be mistaken for good news.

Regulated wrappers may accelerate access, liquidity, legal clarity, and institutional holdability. They can also produce **holders without users**. A spot ETF may hold the underlying asset while shielding its investors from self-custody, private settlement, native fees, proof procurement, and verified-compute consumption. A treasury company may amplify exposure through debt and equity issuance without increasing protocol-native demand proportionally. A daily-reset leveraged wrapper may amplify price further while introducing path dependence and forced rebalancing that has nothing to do with the protocol at all.

Institutional wrappers are not automatically hostile, and the reflex that every ETF is an adversary is as lazy as the reflex that every ETF is adoption. The real question is empirical: **does wrapper growth seed native use, or permanently substitute for it?** Public telemetry must

Wrapper benefit	Monetary cost
Easier legal ownership	Custody centralization
Greater liquidity	Price decoupling from native use
Institutional access	Fee/burn/collateral bypass
Familiar reporting	Loss of privacy and self-sovereignty
More price demand	Potential mechanical amplification

Table 5.2: Institutional holdability is a genuine benefit with a genuine monetary cost. The thesis treats both sides as measurable rather than assuming either.

therefore distinguish accessibility from monetary adoption. [Section 26.3](#) states the holdability requirements; [Section 24.8](#) states the measurements that keep them honest.

These tensions are not defeaters of the thesis; they are **boundary conditions**. The stack does not promise to solve all social and political problems. It promises:

1. Technical guarantees that are real within their scope.
2. Telemetry that makes boundary violations visible.
3. Economic and architectural design that expands the scope of those guarantees over time.

Chapter 6

Layers of the Cypherpunk Stack

The threat model makes one thing clear: this is not a single-layer contest between “good” and “bad” money. It is a contest across an entire stack of machinery, from silicon and power all the way up to governance and law. States under stress will not only try to tax balances or censor individual transactions; they will reach for routing, hardware mandates, identity schemes, update channels, and legal definitions of “safety.” If these services are to support a credible base-asset monetary candidate rather than another platform token, the full stack and holder constituency must survive on all of those planes.

It is useful, therefore, to make the architecture explicit as a set of layers. We already speak of **Layer 0: Verifiable Machines**, which means open hardware and sampled supply chains as base reality. Above that, the thesis has described triad-centric layers in terms of Create/Compute Prove Settle Verify, with Telemetry & Governance threaded through. We can now refine this into a seven-layer stack that stretches from matter to institutions:

- **Layer 0:** Verifiable Machines (and power)
- **Layer 1:** Reachability (communications and transport)
- **Layer 2:** Distribution & Execution (software supply and runtime)
- **Layer 3:** Identity & Claims (who may act, without doxxing)
- **Layer 4:** Truth & Work (proofs, VerifyPrice, and useful compute)
- **Layer 5:** Value & Settlement (privacy rails and non-custodial flow)
- **Layer 6:** Governance & Telemetry (keeping the system honest)

The rest of the document can then be read as a tour through these layers.

Stack Viability $\neq$ Monetary Viability

A working stack is necessary but not sufficient for store-of-value behavior. Parts III–IV show that the stack *can* supply the triad. Part II must still show that a specific asset *captures the value* of that supply. Without the value-capture bridge developed in [Chapter 7](#), the stack is useful infrastructure but not money.

The Duration Warehouse Is Not a Layer Either

Section 3.3 distinguishes duration of the claim from duration of the project. Layers 0–6 specify machines, reach, code, identity, proofs, settlement, and telemetry. None of them is a pension, a directed bank, or a construction-finance book. The missing primitive belongs in the diagnosis of the incumbent order, beside the collateralized sovereign stack (Section 3.2), not as an eighth addition to this ladder. Proofs can audit long-lived credit. They cannot become it.

Layer 0 remains the bedrock. **Verifiable Machines** are the point where cryptography stops being metaphor and touches matter. The chips that sign, prove, randomize, and capture must not be opaque priesthood artifacts; they must be at least partially inspectable and sampled. That means open RTL or microarchitectures where we can get them, open PDKs where politics allow, and structured lot-sampling and imaging where we cannot. It also means treating power and physical plant as part of Layer 0 rather than an afterthought: a prover farm or router that cannot withstand a planned blackout is as brittle as a closed chip. In practice, Layer 0 is “verifiable machines on verifiable power,” which means open designs, measurable side-channel budgets, lot attestations, and micro-grids or backup arrangements that keep the hardware honest and powered when stress arrives.

On top of that sits **Layer 1: Reachability**. If Layer 0 asks “can we trust what the machine does?”, Layer 1 asks “can we talk to it at all?” A repression toolkit that has already learned to pull levers on subsea cables, IXPs, DNS roots, and mobile carriers will aim squarely at transport. Here the goal is simple: make it hard to turn the network off without turning the country off. Concretely, that means encrypted transports that blend into ordinary traffic, obfuscation that defeats naive DPI and active probing, alternate paths through satellite and radio when fiber is throttled, and enough diversity of routes that there is no single “kill switch.” In the stack metaphor, Layer 1 is the oxygen line—if it fails, proofs and privacy become academic. The communications section we introduce later makes this concrete: Tor-class pluggable transports, encrypted handshakes, refraction networking where cooperators exist, plus sat/mesh/sneakernet fallbacks all belong here.

Layer 2: Distribution & Execution answers a quieter but equally potent adversary move: “if we cannot block the packets, we will block the programs.” App stores, CDNs, corporate MDM policies, and automatic update systems are powerful levers. A network whose clients can only be installed or updated with the blessing of a small set of platforms is not neutral; it is merely waiting to be deputized. Layer 2 is where we insist that the code that speaks triad-money can circulate without anyone’s permission. Content-addressed binaries, signed manifests mirrored across jurisdictions, USB/QR/audio installers, and runtimes that can be fetched and verified over whatever transport is available are the practical elements. Execution environments matter too: if the only viable runtime is a locked-down phone OS wired to

one app store and one identity scheme, the repression toolkit has already won. The software distribution section we add later sits here, making update and execution independence a first-class requirement.

If Layer 2 makes sure the code can run, **Layer 3: Identity & Claims** makes sure it can express “who” is acting without collapsing into real-name KYC. The adversary here is the growing web of SIM registration, network identity, and “safety” regimes that treat anonymity itself as suspicious. The stack does not deny that some actions must be accountable; it refuses to equate accountability with global doxxing. At this layer, keys, credentials, and policies are defined in a way that allows entities to prove the right to act (spend, sign, operate a prover) without binding that right to a civil identity unless absolutely necessary. Anonymous or pseudonymous credentials, selective-disclosure proofs, and reputation linked to receipts and behavior rather than to phone numbers belong here. When we later speak of viewing keys, lawful-privacy corridors, and audit-friendly receipts, we are concretely filling in Layer 3.

Only once machines are honest, packets can move, code can run, and actors can be named without being exposed do we reach **Layer 4: Truth & Work**, the core of the original thesis. This is where proofs, VerifyPrice, and useful work live. The job of this layer is to answer, cheaply and publicly, the question “did this actually happen as claimed?” It is here that verification asymmetry is engineered: MatMul-PoUW constructions and ZK proof systems that make checking an order of magnitude cheaper than doing; harnesses that measure p50/p95 verifier time and cost for each workload; and markets that pay only for receipts that pass those checks. In the earlier language of the loop, Layer 4 is the Prove/Verify spine: the circuits, proofs, and reference verifiers that turn “trust me” into “verify me” at scale. Much of the existing text (Verification Asymmetry & VerifyPrice, Proof Factories, AI-PoUW, the modular stack) already describes this layer in detail.

With Layer 4 in place, **Layer 5: Value & Settlement** can safely be money-like. This is where privacy rails, atomic swaps, shielded pools, and cross-chain bridges live. Its task is double: move value without custody or censorship, and do so in a way that preserves privacy by default while leaving a path for auditable disclosure. Here the adversary is both the classical financial repression toolkit (negative real yields enforced through custodians and captives) and the newer chokepoints of KYC’d exchanges, surveilled payment processors, and compliant stablecoin issuers. The response is not rhetoric; it is architectures that keep custody at the edge, settlement neutral in the middle, and content private unless a viewing key is invoked. Adaptor-signature swaps between BTC and privacy assets, lawful-privacy corridors for treasuries, bridge designs that avoid trusted relays, and settlement telemetry (success rates, refund safety, anonymity-set health) are the concrete pieces. The Privacy Rails sections, and the new communications and identity work, all culminate here.

Finally, **Layer 6: Governance & Telemetry** closes the loop. If the lower layers are the ner-

vous system, this is the cortex and immune system: how the network notices drift, resists capture, and adapts under stress. Without this layer, even beautifully engineered proofs and privacy corridors will decay into the usual pattern of oligopolies and black boxes. In practice, Layer 6 is where VerifyPrice observatories, decentralization dashboards, corridor-health monitors, and incident reports live, along with the institutions that respond to them. It is also where operational and legal posture is set: how routers are constrained so they cannot quietly privilege “house” flow; how upgrades are staged so Layer 0 and Layer 4 stay aligned; how incident response works when a corridor breaks or a hardware profile is compromised; how policy engagement happens without creating a single political point of failure. The mantra “no dashboards, no trust” belongs to this layer: if the public cannot see verification cost, neutrality, censorship, and settlement health, the system has already slid back toward platform trust.

Seen this way, the cypherpunk monetary stack is not one clever consensus mechanism but a ladder of seven interlocking layers. Layer 0 and Layer 1 keep the silicon honest and the packets flowing. Layer 2 and Layer 3 ensure that code and identity cannot be quietly turned into chokepoints. Layer 4 turns work into truth with receipts anyone can verify. Layer 5 turns those receipts into private, non-custodial settlement. Layer 6 watches the whole organism and pushes it back toward neutrality when it drifts. The sections that follow simply walk this ladder: beginning with verifiable machines, extending upward into communications resilience, distribution, and identity, and then re-entering the territory already mapped (proofs, compute, settlement, governance) with a clearer sense of what each layer must survive.

6.1 Modules in Outline: Applications and Primitives

The layered picture tells us **where** things must live; the modular picture tells builders **what** they actually touch. We will use four reference applications as running examples throughout the rest of the thesis:

1. **Private treasury & payroll:** Pays staff and vendors non-custodially over privacy rails, emits receipts that auditors can check, and treats “who got paid what” as a matter of viewing keys, not public gossip.
2. **Media provenance & authenticity:** Cameras, sensors, and editors emit cryptographic lineage that survives platform stripping; payments to creators and data providers are conditioned on the presence of such receipts.
3. **Verified inference:** A market where model owners and service providers sell inferences that come with succinct proofs or hybrid proofs-of-logits, all priced and measured via VerifyPrice.
4. **Proof/compute procurement:** A rail where developers and treasuries can buy standardized units of “proofs” and “verified FLOPs” as futures or spot capacity and treat them as balance-sheet items next to BTC.

To support these, we rely on a small set of reusable primitives:

- A **Proofs-as-a-Library SDK (PAL)** that lets developers declare claims (“prove this computation,” “prove this provenance”) and compiles them to multiple proving backends and useful-work miners.
- A **Privacy Rails Kit (PRK)** that expresses pay-for-proof intents and executes refund-safe settlement over BTC↔ZEC/XMR corridors or shielded pools.
- A **Proof Interface Definition Language (PIDL)** that defines the minimal receipt that binds a claim, a proof, a workload ID, an SLA tier, and timestamps into a portable object.
- A **VerifyPrice observatory** that provides public metrics for how cheap verification actually is for each canonical workload.
- A set of **market and telemetry primitives** (SLA escrows, neutral routers, bridge-safety templates) that make it possible to treat proofs, privacy capacity, and verified FLOPs as commodities rather than favours.

Later sections turn this into a full “substrate kit” of twelve primitives and four reference applications. Until then, we will use the names PAL, PIDL, privacy-rails kit, proof factories, and VerifyPrice observatory as shorthand for these building blocks.

Part II

The Triad and the Monetary Candidate

Part I argued that the old guarantees are failing under debt arithmetic, AI, and platform control, and that a seven-layer cypherpunk stack is needed to survive repression. This Part looks at the same system from the **monetary angle**.

The claim is conditional: **Privacy, Proofs, and Compute can support a bearer monetary candidate only if their full service path remains stress-deliverable, value accrues without bypass, and a persistent self-custodied constituency bears the asset's loss.** The services do not become money by being necessary, and they need not share one token. This Part separates the base asset from typed Work Credits, evidence artifacts, project credit, and derivatives.

Chapter 7

The Triad and the Monetary Candidate

Traditional monetary regimes pick a **base reality** and build promises on top of it:

- Gold: geology + metallurgy.
- Fiat: law, taxation, and war.
- Bitcoin: thermodynamics and code.

Each regime implicitly answers two questions:

1. *What counts as real work?*
2. *What object or capacity will we treat as the canonical memory of that work?*

In a civilization where intelligence is largely machine-executed and verification can be cheap and public, the answer shifts: **work is what machines can do and humans can verify cheaply**, and the “object” that records that work no longer has to be a metal or a pure ledger entry. It can be a *capacity*.

This thesis treats three capacities as monetary primitives:

- **Privacy**: censorship-resistant settlement that preserves agency.
- **Proofs**: portable attestations of computation and provenance.
- **Compute**: useful work wrapped in succinct guarantees.

Each is:

- **Indispensable** in a dense digital economy (no safe commerce without privacy; no safe coordination without proofs; no AI without compute).
- **Verifiably scarce** at any point in time (bandwidth, cycles, proof capacity are bounded by physics and capital).
- **Cheap to verify** (you can check whether you have privacy, a valid proof, or a verified FLOP without trusting a platform).

The physical and service base in this frame is the **networked capacity to deliver Privacy, Proofs, and Compute under adversarial conditions**. It is not itself money. Those services may support an associated bearer base asset as a conditional monetary candidate only if all nine links pass; Work Credits remain typed claims on service or capacity.

A few distinctions help:

Services vs. instruments: DVC measures how much verifiable Privacy/Proofs/Compute per unit time the stack can deliver at target SLOs. Work Credits slice specified service capacity into transferable claims; staking and LP positions are derivatives or operating claims; the separate

base asset is the only monetary candidate.

Nominal vs. real: Nominal pricing of the triad will oscillate in terms of fiat and BTC. Real value is “does this capacity still buy me censorship-resistant settlement, proofs, and verified FLOPs when repression and AI get worse?”

Symbol vs. utility: Gold and BTC partly trade as symbols. The triad trades as **plumbing**: “can I still pay people, prove things, and run intelligence without asking permission?”

For the associated base asset, recurring service demand is necessary but insufficient:

- Treasuries and individuals purchase **private settlement** to escape surveillance and yield-curve control.
- Platforms, enterprises, and states purchase **proofs** to secure provenance, compliance, and audit trails in an AI-polluted information environment.
- AI labs, agents, and applications purchase **verified compute** to sell trustworthy services.

The claim is not that the triad services behave like a reserve asset. “Private Money” and “AI Money” are analytical lenses on settlement and compute demand. Monetary premium, if any, belongs to the associated base asset after deliverability, non-bypassability, holder quality, and the remaining chain conditions pass.

7.1 The Conditional Chain, End to End

Chapter 1 promised that the thesis would establish its claim through a conditional chain rather than an assertion. That chain is defended link by link across six Parts, which makes it easy to lose. This section walks it once, in one place, and says where each link is argued and what would sever it.

Utility demand → standardized work → receipts → stress-deliverable service → markets → fee flows → a scarce asset (only under non-bypassable accrual) → a persistent, self-custodied, loss-bearing, regime-responsive holder constituency → a possible store-of-value premium.

Every arrow is a conditional. None is asserted as inevitable, and the chain is only as strong as its weakest link—which is why we name the weak ones explicitly at the end rather than leaving the reader to find them.

1. **Utility demand exists and persists.** A dense, synthetic-media, AI-mediated economy must keep buying private settlement, portable attestations, and verified compute through every cycle—including, and especially, under financial repression (Chapter 3, Chapter 8, Chapter 9, Chapter 10).

Severed if: the demand turns out to be cyclical discretionary spend rather than structural necessity.

2. **Demand standardizes into canonical work.** Diffuse demand becomes an economic unit only when workloads are specified precisely enough that two providers are selling the same thing—canonical workload registries, hardware profiles, and SLA tiers (Section 20.3, Section 20.4).
Severed if: workloads stay bespoke, so no fungible unit forms and there is nothing to price.
3. **Work produces cheap-to-check receipts.** Standardized work emits portable artifacts—PIDL receipts binding claim, proof, workload ID, tier, and timestamps—and verification stays far cheaper than production, $r(W) = v(W)/p(W) \ll 1$ (Section 20.2, Section 20.5, Appendix A).
Severed if: verification cost creeps toward production cost, at which point “anyone can verify” collapses into “trust the prover.” This is Red Line 1, the hinge (Section 28.6).
4. **Receipts represent stress-deliverable service.** A receipt is evidence, not capacity. For each disruption scenario, Delivered Verified Capacity traces the full path from energy and fuel through firm power, switchgear, cooling, hardware, network, workload, proof, verification, settlement, and usable service; issuance is bounded by the surviving max flow and active minimum cut (Section 15.9).
Severed if: gross installed capacity or benign-state output materially exceeds the service that survives substitution latency and common-cause failure.
5. **Receipts enable markets.** Cheaply checkable receipts over deliverable service let strangers transact without trusting each other: proof factories, neutral routers, SLA escrow and slashing, and procurement of proofs and verified FLOPs as spot or forward capacity (Section 20.10, Section 22.1, Section 22.6).
Severed if: routers, provers, or matching engines capture the market and rent replaces work—Red Line 3 territory.
6. **Markets produce fee flows.** Working markets generate recurring, measurable revenue rather than one-off grants or subsidies, visible as fee-coverage ratios against the security budget (Section 20.10, Section 24.9).
Severed if: fee coverage collapses and workload demand proves to be speculative rather than budgeted.
7. **Fee flows accrue to a scarce asset—only under non-bypassable rules.** This is the load-bearing link and the one most theses skip. Utility does not become money by being useful. The [Value Capture Lemma](#) states the five conditions: required fee medium, meaningful burn or retirement, staked collateral, disciplined issuance, and no bypass channel delivering equivalent service without touching the asset (Lemma 11.1, Section 11.11, Section 11.12).
Severed if: users obtain equivalent triad capacity through hyperscalers, stablecoins, or custodial APIs without the asset—Red Line 6 (Section 28.6, Section 26.2).
8. **A monetary-risk warehouse forms.** Gross native demand is not an anchor. A persistent, self-custodied, loss-bearing, regime-responsive constituency must retain balances, lengthen holding periods, and accumulate countercyclically rather than merely acquire fees just in

time, recycle operator inventory, satisfy collateral rules, or hold wrappers (Section 24.8.8, Section 24.8.7).

Severed if: native demand disappears into immediate sell-through, leverage, wrappers, or forced deleveraging during the very stress in which the asset is meant to preserve agency.

9. **A scarce, captured, anchored asset may earn a store-of-value premium.** If the asset also remains liquid, neutral, verifiable, legally holdable, and agency-preserving, it can plausibly earn monetary premium (Chapter 32, Chapter 14, Chapter 4). This final link, not fee accrual or gross demand, is where the premium is argued to live (Section 11.10.8).

Severed if: any of the nine store-of-value requirements fails, or telemetry showing them is captured (Red Line 4).

The Chain Says Nothing About Price

Completing all nine links would establish only that the base asset is a defensible monetary candidate. It would not establish when, or with what volatility, price reflects that. Wrappers, leverage, dealer hedging, and allocation mandates can move price with no protocol use at all, in either direction (Corollary 11.1). The chain is measured on the protocol, capacity, and native-buyer boards; price is interpreted separately through VerifyFlow (Section 24.8) and never cited as evidence for the chain.

The chain rests on something beneath it. Links 2 through 6 consume bulk power, and link 3—the verification hinge—is denominated in real resources on reference hardware. That makes the chain dependent on conditions that states partly set: energy availability and pricing, the obtainability of unprivileged reference hardware, and the continuity of every conversion edge through usable service (Section 5.5, Section 15.10.1, Section 15.9). Sovereign Optionality measures pathway diversity; Delivered Verified Capacity measures surviving flow. The chain is not suspended in mathematics; it is bolted to a grid and held on balance sheets.

Where the chain is weakest. Honest accounting rather than uniform confidence:

The medium ratings are the point. Non-bypassability is a magnitude: bypass is never impossible, only priced, and the marginal buyer's willingness to pay for the protocol's differential properties has not been estimated. Delivered service is not nameplate capacity: substitution latency and common-cause cuts can erase apparently diverse supply. Holder demand is not risk absorption: fee purchases, burns, collateral, and wrappers can coexist with no constituency willing to own the asset through a drawdown. The premium, if it exists, comes from link 9 together with the state-contingency of differential value (Section 11.10.8). Ratings that go up are self-congratulation; ratings that identify the unmeasured bridges carry information.

The rest of this Part defends links 7 through 9, where the monetary argument actually lives.

Link	Strength	Why
Utility demand	Strong	Structural drivers are already observable
Standardized work	Medium	PoUW constructions are a research direction, not a shipping monetary base
Receipts	Strong	Structurally well specified; VERIFYPRICE makes it measurable
Stress-deliverable service	Medium	Scenario max-flow/min-cut and common-cause data are specified but not yet demonstrated at scale
Markets	Medium	Market microstructure for proof and compute capacity is the least developed part of the design
Fee flows to asset	Strong	Formalized as a lemma with public boards—but establishes <i>cash-flow accrual</i> , not monetary premium (Section 11.10.7)
Non-bypassability	Medium	Written as a binary; the economics requires a magnitude (Δ) that is nowhere estimated (Section 11.10.4)
Native holder anchor	Medium	Transactional, collateral, and wrapper demand do not establish countercyclical loss-bearing capacity
SoV premium	Medium (conditional)	Nine requirements, fifteen red lines; premium derives from link 9 plus state-contingency, not from fee accrual (Section 11.10.8)

Table 7.1: Link-by-link strength of the conditional chain. Six links remain Medium or Medium (conditional): standardized work and market formation are engineering-maturity risks; stress-deliverable service is a physical-topology and common-cause risk; non-bypassability is an unmeasured magnitude; the native holder anchor is an unproven risk warehouse; and monetary premium remains conditional on the entire chain.

Parts III and IV defend links 2 through 6 by building and measuring the stack that supplies them. Parts V and VI make every link falsifiable.

7.2 Topological Scarcity and the Pressure–Capacity Corridor

Topological Scarcity Lemma

A nominal stock of resources, capital, or capacity does not secure a monetary service. The relevant quantity is the stress-deliverable flow through the narrowest non-substitutable edge connecting input to holder utility.

Aggregate abundance and local scarcity can therefore coexist without contradiction. Crude is not diesel when refining is the minimum cut; savings are not duration absorption when no balance sheet can hold DV01; installed accelerators are not verified service when power, cooling, firmware, connectivity, verification, or settlement fails. The topology—the transformation path, location, actor, and horizon—determines the service.

The physical-conversion example is adapted selectively from Doomberg’s *Endangered Specious*; the duration-absorption example from Michael Green’s *Anchors Aweigh, My Boys*. Both are analyst commentary used for mechanism rather than quantities. The lemma and its application to verified service are ours.

Let R denote regime pressure, $\Delta(R)$ the holder’s differential value over the best institutional substitute, $A_{\text{full-stack}}(R)$ the availability of the complete service path under that pressure, and $H(R)$ the holder’s practical ability to self-custody, transact, prove, and exit. The monetary service scales with the product of the three, on the product-form premise

$$\Pi_{\text{monetary}}(R) \propto \Delta(R) \times A_{\text{full-stack}}(R) \times H(R).$$

This is the **Pressure–Capacity Corridor**. Its likely shape is an inverted U, not a crisis escalator. At low pressure, courts, custodians, and indemnitors work, so Δ is thin. At moderate pressure, substitutes weaken while the open stack and holder remain operational, so differential value can rise. At extreme physical or political pressure, power, networks, hardware, liquidity, custody, or exit may fail and the service falls with them. More crisis does not imply more premium.

7.3 The Triad Coherence Test

Privacy settlement, portable proof, and verified compute are different goods with different users, bottlenecks, collateral needs, political risks, and duration warehouses. A common asset

is justified only if shared security, settlement, liquidity, and transaction-cost benefits exceed cross-subsidy, governance, and wrong-way-risk costs.

Architecture	Strength	Burden of proof
A: One native monetary asset	Strongest direct value capture and shared liquidity	Greatest reflexivity, governance scope, and cross-domain collateral risk
B: Neutral reserve plus service-specific credits	Clean separation between reserve money and typed service claims	Weaker new-token capture; requires interoperable settlement and independent credit markets
C: Shared settlement plus modular domain collateral	Common payment rail with service-specific risk containment	More operational complexity; collateral fragmentation and bridge governance must remain legible

Table 7.2: Triad architectures under test; none is pre-selected.

For each architecture ask: Do the services share a security budget and users? Does demand for one improve the economics of the others? Are bottlenecks complementary or correlated? Does common settlement reduce transaction costs more than it creates cross-subsidy? Can one constitution credibly govern all three? Does one service dominate issuance, fees, or political risk? Architecture A is a candidate, not the thesis.

7.4 Monetary Objects and Value Capture

The thesis claims that triad capacity can earn a “durable store-of-value premium.” For that claim to be testable, we must answer four questions without poetry:

1. **What exactly is the asset?**
2. **What do users pay in?**
3. **How does holding capture value?**
4. **Why doesn’t value leak entirely to operators?**

Question 1: What is the asset? The triad stack can issue several kinds of holdable instruments:

The **base asset** is only a conditional monetary candidate: the unit in which fees may be paid, burns may occur, and collateral may be posted. Work Credits can be valuable, transferable, workload-specific, location-specific, SLA-specific, and temporally perishable without becoming savings instruments. FCRs, FERs, and PIDL receipts are evidence; project notes

Instrument	What It Represents	Scarce?	Transfer?	SoV?
Base Asset	Conditional bearer monetary candidate; fee and settlement unit	Yes	Yes	Candidate only
Work Credits	Typed capacity or service claims, bounded by deliverable output	Yes	Yes	Not presumed
FCR/FER/PIDL artifacts	Evidence of physical conditions or completed work	No	Yes	No
Project Notes	Explicit duration-bearing infrastructure credit	Yes	Yes	No
LP/Staking	Rights to fee share from corridors, pools, validators	Yes	Sometimes	Derivative
Capacity Vouchers	Prepaid access at fixed rates	No	Yes	No

Table 7.3: Instrument hierarchy. Utility, evidence, credit, derivatives, and the conditional monetary candidate remain distinct.

warehouse duration and bear default risk; LP and staking positions are derivatives or operating claims.

A terminology contract. This thesis uses two distinct words that are easily confused, and the distinction carries real analytical weight:

- **Native instruments** are protocol-internal claims: the base asset, typed Work Credits, staking and LP positions, corridor claims, capacity vouchers, and evidence artifacts. They touch the protocol, but only the base asset is evaluated as a monetary candidate.
- **External financial wrappers** are conventional market products written *on* the asset: spot ETFs, exchange-traded products, treasury companies, custodial balances, futures, options, swaps, and leveraged or inverse ETPs. They do not touch the protocol. Holding them exercises nothing.

Both are sometimes called “wrappers” in casual usage. From here forward, “wrapper” without qualification means an *external financial wrapper*. The full object hierarchy runs from physical capacity out to recursive financial claims, and monetary treatment differs sharply along it.

How to read Questions 2–4. The answers below are stated in the reference-design register: they specify *routing rules* — where fees go, what must be locked, what may not be inflated away. They are design commitments, not findings, and two things they deliberately do not do. They do not establish that the routed fee is *large*: that is the level question, taken up in [Section 11.10](#), where the sustainable fee is bounded by differential value Δ rather than by any-

Object layer	Example	Touches protocol?	Path dependence	Monetary treatment
Service capacity	Privacy, proof, compute throughput	Yes	No	Non-monetary input measured by DVC
Typed service claim	Work Credit / capacity voucher	Yes	Medium	Not presumed money or SoV
Native monetary object	Base asset	Yes	Low	Conditional SoV candidate
Native derivative	Staking or LP shares	Yes	Medium	Secondary claim
Custodial claim	Exchange balance, spot ETF share	Indirectly	Low–medium	Price exposure, not full monetary function
Synthetic wrapper	Futures, swaps, options	No	Medium–high	Financial exposure only
Daily-reset leveraged wrapper	2× / 3× ETP	No	Extreme	Trading product; explicitly not SoV
Recursive wrapper	Leveraged product on another wrapper	No	Extreme, nonlinear	Systemic-risk instrument

Table 7.4: Object hierarchy from base capacity to recursive financial claims. Monetary properties do not survive the descent: a daily-reset leveraged wrapper cannot inherit the underlying asset’s store-of-value status, because its return depends on the path, not merely the endpoint (Appendix H.9).

thing in this section. And they do not establish that any of this is monetary premium rather than a cash-flow claim: [Section 11.10.7](#) withdraws exactly that reading, and [Section 11.10.8](#) re-locates the monetary claim to a holder-side service flow. The percentages below are illustrative of a design space, not measurements of a market.

Question 2: What do users pay in? All triad services are priced in the **base token**:

- **Proofs:** Pay base tokens to provers; portion burned.
- **Privacy settlement:** Pay base tokens for corridor fees; portion burned.
- **Verified compute:** Pay base tokens for inference/MatMul; portion burned.
- **Staking/collateral:** Provers, routers, and LPs must lock base tokens to participate.

This creates **structural demand**: every use of the triad requires acquiring the base token.

Question 3: How does holding capture value? Value accrues to holders through four channels:

1. **Fee burns** (30–50% of fees): Deflationary pressure proportional to usage.
2. **Staker yield** (20–40% of fees): Income for holders who stake.
3. **Operator share** (20–40% of fees): Incentive for provers/routers/LPs.
4. **Collateral requirements** (10–20% of capacity value): Immobilizes float and sets a valuation floor; [Section 11.10.9](#) shows why it is a floor and a slope effect rather than proportional accrual, and bounds the fixed dollar quantity it can lock.

The first three channels are a **partition of one fee pool**, not three independent receipts: burn share, staker yield, and operator share must sum to 100% of every fee collected. The bands above are alternative interior choices within that partition, not independent ranges, and any point allocation — such as Design A’s 40/30/30 split ([Section 7.5](#)) — must close the identity. Reading the bands as additive would sum the maxima to 130% of the pool, which is not an allocation the design can make.

Question 4: Why doesn’t value leak entirely to operators? The “utility token trap” occurs when operators capture all economic value while token holders merely provide exit liquidity. The stack avoids this through:

- **Fee burns:** 30–50% of fees are permanently destroyed, benefiting all holders.
- **Staking yields:** Passive holders can stake to earn fee share without operating infrastructure.
- **Collateral requirements:** Operators must hold significant tokens, aligning their interests with holders.
- **Governance rights:** Token holders vote on fee splits, issuance changes, and protocol upgrades.
- **Issuance constraints:** Base-asset issuance follows the constitutional schedule and is independent of Work Credit dynamics; it is *Work Credits* that are minted only against verified

capacity growth, with minting beyond real capacity triggering SLO breaches visible in dashboards.

Net effect, routing only: these rules route a fee of size $f \cdot Q$ to retirement, staking, and operator share. Whether f can be positive at all is a function of differential value Δ , not of routing; whether the routed revenue constitutes monetary premium is a separate question this Part takes up in [Section 11.10](#) and answers in the negative at [Section 11.10.7](#).

7.5 Three Reference Designs

To make the thesis testable, we commit to three concrete designs. Implementations may vary, but at least one must be viable for the SoV claim to hold.

Design A: Base Token as SoV (Primary)

- **Issuance:** Fixed schedule with halvings (like BTC) or capacity-linked ceiling.
- **Fee medium:** All triad services priced in base token.
- **Burns:** 40% of fees permanently destroyed (a point choice inside the 30–50% reference band).
- **Staking:** 30% of fees to stakers; 10–15% collateral requirement.
- **Operators:** 30% of fees to provers/routers/LPs.
- **Governance:** Token-weighted voting on parameters.

SoV properties: Credible scarcity (capped + burns), native demand (fees), duration-neutral (no coupons), cheap verification (VerifyPrice dashboards).

Risk: If demand stalls, burns decline and scarcity weakens.

Design B: Work Credits as Capacity Vouchers (Hedge Instrument)

- **Issuance:** Minted against verified work (FERs + proofs); no fixed cap.
- **Redemption:** Burnable for priority access to proofs/compute/settlement at SLA-guaranteed rates.
- **Expiry:** Credits decay or expire after N years to prevent hoarding and bank-run dynamics.
- **Transferable:** Yes, but primarily used for cost hedging, not long-term savings.

Use case: Enterprises hedge against compute cost spikes; treasuries lock in future settlement capacity.

Not a SoV: Supply expands with capacity; expiry prevents indefinite accumulation. This is infrastructure hedging, not a store of value.

Design C: Triad Index Token (SoV Candidate)

- **Backing:** Diversified revenue streams across privacy corridors, proof pools, and compute markets.
- **Value capture:** Protocol revenue used for periodic buyback-and-burn or dividend distribution.
- **Issuance:** Fixed supply; no new minting after genesis.
- **Governance:** Index holders vote on revenue allocation and rebalancing.

SoV properties: Diversified exposure to triad demand; fixed supply; yield via buybacks or dividends.

Risk: Concentration in specific corridors or proof markets; governance capture.

Which design does the thesis endorse? None by default. These instrument designs sit inside the broader architecture comparison of [Section 7.3](#). A fixed-supply base token may be tested as the monetary candidate; capacity vouchers remain service hedges; an index token is a revenue-bearing security whose buybacks or dividends do not establish moneyness. The text uses “base asset” for the monetary candidate and “Work Credit” only for a typed capacity or service claim.

What actually differs, and who loses. The three designs are not three shades of the same bet; each one moves the risk onto a different balance sheet. Design A concentrates it in the base asset: if demand stalls, burns decline, and the scarcity story thins with it — the holder discovers that a burn schedule is a claim on volume, not a property of the token. Design B pushes the risk onto the service buyer: capacity vouchers hedge compute cost beautifully until capacity is scarce in exactly the state the buyer hedged for, at which point delivery terms rather than price are the binding constraint, and the hedge is only as good as the stressed path behind it. Design C hands the risk to governance and to the revenue mix: an index token is a claim on whatever the protocol happens to earn, so concentration in particular corridors or proof markets is not a footnote to the thesis but the thesis, and buybacks do not convert a revenue stream into money ([Section 11.10.7](#)). Watch the same three numbers in all three cases — fee share actually paid in the native asset, net supply after burns, and the share of exposure sitting in wrappers ([Section 11.13](#)) — because [Section 11.10](#) takes up the question these boxes deliberately do not answer: not where the fee is routed, but whether there is a fee to route.

Chapter 8

Privacy as Private Money

Cash used to be **default private money**: anonymous, bearer, final on receipt. In a world of KYC'd banks, programmable payments, and networked surveillance, that role has decayed. At the same time, the repression playbook (negative real rates + capital controls) makes **bearer-like savings** economically necessary, not ideologically optional.

“Privacy” in this thesis is not romantic opacity; it is **the ability to hold and move value without chokepoints, plus the option to disclose on your own terms.**

Interiority has a history. The capacity this chapter argues for has a five-hundred-year arc behind it, and the arc is worth naming because it reframes what is being defended. The private interior self is not a natural constant that technology is eroding; it is, on the account of the orality-literacy literature (Ong 1982; Sources), substantially a *product* of a technology: sustained silent reading, a practice individuals had to acquire, that only became mass practice with print. Interiority was an externality of a medium, and it has since been mistaken for the default state of persons. On that reading, surveillance does not violate a natural privacy so much as withdraw a subsidy that a particular medium had been paying, quietly, for centuries. The implication for this chapter is not that the case for private settlement strengthens — the repression economics stand on their own — but that the loss is larger than it looks: what is being unbuilt is not an amenity but a specific, technology-dependent form of the person that took centuries to build and that no one planned. Defending it is therefore also not an amenity.

Concretely:

Bearer-like holding: Keys, not accounts, define control. Custodians may exist, but custody is an *optional service*, not a mandatory chokepoint.

Non-custodial settlement: Cross-asset flows (e.g., BTC↔ZEC/XMR) execute as atomic swaps or corridor protocols; neither side needs to trust a centralized intermediary.

Auditable by consent: Viewing keys and structured receipts allow specific flows to be disclosed to auditors without exposing the entire graph.

When these properties hold, **private settlement capacity itself** starts to behave like a monetary asset:

- A treasurer facing capital controls is not just asking “what’s the yield?” but rather, “*can I still get value to my people next year if on-/off-ramps are throttled?*”

- A dissident journalist or NGO cares less about upside and more about *“will this still be here and spendable if my local banks freeze?”*

Blockchain may be waiting for its SSL moment. The analogy is suggestive: in 1994, putting credit card information on the Web seemed reckless until SSL made encrypted commerce viable. E-commerce grew into a multi-trillion-dollar industry once confidentiality became default. Today, most public blockchains are “public by default” the way HTTP was—every transaction visible, every address linkable. Privacy is plausibly the bottleneck for mass institutional adoption.

The institutional version is straightforward: no enterprise wants payroll, vendor rates, or treasury operations visible by default. The unlock is one-click, non-custodial BTC↔XMR/ZEC with refund-safe UX and clear settlement analytics—privacy as a product, not a promise.

In that environment:

- The **rails** (privacy corridors, shielded pools) earn fees for providing unseizable, auditable settlement capacity.
- The **claims** on those rails (e.g., corridor LP positions and Work Credits) are service, operating, or derivative claims. They may be held, but are not presumed stores of value.

You can think of **Private Money** as:

“Rights to future, censorship-resistant, auditable settlement capacity.”

The SoV properties from [Chapter 4](#) map naturally:

- **Credible scarcity.** Capacity is constrained by bandwidth, cryptographic verification costs, and capital at risk.
- **Cheap public verification.** Anyone can verify that a transaction was correctly formed, swapped, or refunded.
- **Censorship-resistance & portability.** Flows ride over non-custodial corridors; exit options span multiple assets and jurisdictions.
- **Neutrality & permissionlessness.** Adversarially diverse relays, routers, and LPs; public metrics on concentration.
- **Native demand.** Demand is created by repression itself.
- **Lawful privacy.** Viewing keys + PIDL receipts mean regulated entities can prove compliance without deanonymizing entire networks.
- **Duration-neutrality.** Claims participate in fee flows and scarcity premiums, not fixed coupons.
- **Asset-level value capture.** Fees, collateral, and settlement are denominated in the base asset, so consuming the service touches the asset rather than bypassing it ([Section 11.12](#)).
- **Agency preservation.** Selective disclosure and non-custodial settlement equip the holder rather than managing them ([Item 9](#)).

The ninth requirement carries a capability reading the eighth does not, and the distinction matters because a ledger can satisfy the formal properties while failing the practical one. Formal access — a valid balance, a working key, a legal right to transact — is not the same as capability: command over the necessities of an independent life. A perfectly accurate ledger can faithfully preserve an unequal starting state; making claims harder to alter does not broaden who can acquire them (“The Vibecection Was Real”; Sources). The requirement as enforced (Section 24.9) therefore reads on both layers: the Agency Preservation Board tracks non-custodial usage share and disclosure scope — the formal layer — and the participation telemetry of Section 3.6 tracks whether users can in practice acquire and hold meaningful claims through wages rather than inheritance or incumbency — the capability layer. A stack that passes the first and fails the second has hardened the incumbent distribution more efficiently than fiat did.

On a balance sheet, Private Money can coexist with BTC and fiat, but it plays a different role:

- **BTC:** thermodynamic base, global risk asset, macro hedge.
- **Fiat:** near-term unit of account, legal tender, credit medium.
- **Private Money:** repression-hedged rail, a way to ensure you can still pay and be paid without being fully seen.

Chapter 9

Proofs as Attestation Money

If Privacy is the hull, **Proofs** are the **receipts**.

Proofs Are Not Truth

Proofs are bounded attestations, not truth. A proof can show that a computation followed a circuit, that a file descends from a signed provenance chain, or that a settlement rule executed correctly. It cannot prove that a camera was pointed at what matters, that an event was representative, that a witness is honest, or that an AI output is substantively correct. Throughout this chapter, “proof” means “cryptographic attestation under stated assumptions,” not semantic or social truth.

In a world of synthetic media, platform moderation, and “liar’s dividend,” the scarce thing is no longer *content* but **trustworthy provenance and computation history**. Proofs are the mechanism layer that answers, cheaply and publicly:

“Did this actually come from where it claims, and did the computation actually run as stated?”

Clarification: Proofs vs. Proof-Backed Instruments Proofs themselves are *infinitely replicable outputs* once generated. A ZK proof can be copied and verified anywhere. On its own, a proof is not a scarce bearer asset—it is an attestation.

What *is* scarce is:

- **The ability to produce valid proofs at scale** (requires compute, hardware, energy).
- **Rights embedded in instruments** that reference proofs (Work Credits, staking positions, capacity claims).

So we distinguish three roles proofs play in the monetary stack:

1. **Proofs as a commodity market:** Standardized attestations priced by VerifyPrice. You buy proofs the way you buy bandwidth—as a priced input to operations.
2. **Proofs as collateral enablers:** Proofs make *other* contracts collateralizable. An inference SLA backed by proofs of correct execution can be used as collateral because the proof makes default detectable.
3. **Proof-backed service instruments:** Work Credits tied to verified proof workloads are typed service or capacity claims whose integrity is enforced by proofs. They are not presumed

money.

“Attestation Money” survives as an analytical lens or legacy label for proof-demand economics. The instruments beneath that lens are proof-backed service claims; any monetary candidacy belongs to the separate base asset.

From a monetary perspective, two features matter:

1. **Proofs travel:** A PIDL receipt can be moved across systems, archived, or collateralized. It outlives any single platform’s UX, TOS, or reputation.
2. **Proofs can be priced and standardized:** Once you know the VerifyPrice of a workload, you can treat “valid proof of workload *W*” as a commodity unit.

This gives rise to an **Attestation Money analytical lens**:

- Media platforms, insurers, and courts demand proofs of origin, editing history, and custody.
- AI services demand proofs that a model of a given hash ran on given inputs with given bounds.
- Enterprises demand proofs that compliance computations actually ran.

Initially, proofs are purchased as **opex** (“we pay per proof”). Over time, markets will create **claims on future proof capacity**:

- Reservations or futures on proof-of-provenance capacity for a media network.
- Proof pool shares that entitle holders to a portion of fees from high-value workloads.
- Work Credits minted against proof workloads that meet certain VerifyPrice and SLO thresholds.

Under the legacy **Proof/Attestation Money** analytical lens, these instruments are typed claims on proof capacity, not monetary instruments. Work Credits remain service claims; only the separate base asset may become a monetary candidate if the full conditional chain passes:

“Rights to future, standardized, verifiable attestations about data and computation.”

From the triad perspective:

- Privacy protects *who* is involved.
- Proofs protect *what* actually happened.
- Compute powers *how* we arrive at outputs.

The “Attestation Money” lens describes demand arising where **verification, not vibe, mediates trust**. Claims on proof capacity remain service instruments; they are not what this thesis asks holders to treat as money.

Chapter 10

Compute Through the “AI Money” Lens

Compute has always been an economic input: first as muscle, then as steam and electricity, now as FLOPs. The AI boom made this explicit: GPUs, TPUs, and datacenters are priced like oil fields.

Raw compute is not even a standardized service claim:

- Most compute markets are opaque capacity rentals (cloud contracts, colo deals).
- There is no standardized unit of **verified work**; only hours and instance types.
- There is no cheap, public way to verify that a claimed computation actually ran correctly.

Research Maturity Caveat

Verified inference remains technically immature. Only Tier A cryptographic verification (full ZKML) should support high-assurance capacity issuance without an additional trust haircut. Probabilistic or TEE-backed inference verification (Tiers B and C) receives haircuts or lower-assurance service status. The service-value claim weakens if verified compute trades at no durable premium to ordinary cloud compute; that result would say nothing by itself about monetary premium. See [Chapter 22](#) for the tier taxonomy.

The “AI Money” analytical lens reframes verified-compute demand by insisting on:

- **Canonical workloads:** Matrix multiplications, FFTs, and core model primitives.
- **Proofs of useful work (PoUW):** Miners/provers earn rewards for producing proofs that these workloads ran correctly, with verification asymmetry.
- **Verification economics:** VerifyPrice turns “one unit of verified MatMul at dimension n ” into something that can be priced and traded.

In that context, **verified compute capacity** becomes a standardized service commodity:

“Rights to future, standardized, cheaply verifiable units of useful compute (verified FLOPs).”

10.1 Compute Deflation and What Remains Scarce

A skeptical reader will object: “Compute gets cheaper every year. Moore’s Law drives raw FLOP cost down. How can rights to compute retain service value when the underlying commodity deflates?”

This is a serious objection. Here is the response:

What deflates:

- Raw FLOPs per dollar (secular decline).
- Cost of unverified, permissioned compute from hyperscalers.
- Commodity inference for non-sensitive workloads.

What remains scarce:

- **Verified FLOPs with receipts:** Compute where correctness is cryptographically proven, not vendor-asserted.
- **Policy-constrained capacity:** Compute that runs under specific jurisdictional, privacy, or compliance constraints.
- **Censorship-resistant access:** Compute that cannot be denied by TOS changes, sanctions, or platform bans.
- **Priority access under congestion:** During demand spikes, priority slots are scarce even if raw capacity is abundant.
- **Specific hardware profiles:** Compute on L0-C or L0-D grade hardware may remain scarce even as closed alternatives proliferate.

The instrument should reference capacity share, not “one timeless FLOP”: Under the “AI Money” lens, the relevant service claim is not “rights to 1 FLOP forever.” It is:

“Rights to X% of verified compute throughput under SLA Y on hardware profile Z.”

This is analogous to how oil futures reference “barrels of WTI crude delivered at Cushing”—not “energy” in the abstract. The specificity (verification, SLA, profile) is what creates persistent scarcity.

Telemetry that detects deflation risk: If raw compute deflation outpaces the scarcity premium of verification/censorship-resistance, the following metrics will signal it:

- VerifyPrice for verified FLOPs converges toward unverified market rates.
- Capacity utilization on verified networks falls below thresholds.
- Fee revenue from compute workloads declines as a share of total fees.

These are observable. The thesis predicts that a verification service premium may persist because demand for trustworthy AI is structural and the compliance/sovereignty premium may grow as AI becomes more consequential. If telemetry shows otherwise, the compute service case weakens. Neither outcome determines whether the separate base asset earns monetary premium.

Verified-compute service instruments can include:

- **Work Credits:** typed claims minted against verified-compute contributions and redeemable or transferable for specified future workloads.
- **Compute futures:** for specific workloads (e.g., “ X verified inferences at model M with SLO Y ”).
- **Stake in proof factories:** whose entire business is converting energy + silicon into verified FLOPs with transparent VerifyPrice.

Why can these be valuable service claims rather than arbitrary utility tokens? Because:

- **Demand is global and secular:** As long as AI is an input to value creation, there is demand for FLOPs.
- **Verification is public:** Anyone can check that a Work Credit corresponds to a valid proof and a stated DVC envelope.
- **Terms are explicit:** Workload, tier, delivery window, redemption, and expiry determine value rather than an implied monetary promise.

Chapter 11

Work Credits: Energy-Anchored Claims

Every era that mechanizes work eventually writes claims on it. Grain receipts became banking; warehouse receipts became money-market funds; and in each case the instrument was born useful and lived or died on one question — did the claim stay tethered to the thing in the warehouse? The tether is everything. A receipt that outruns its warehouse is a promissory note wearing a receipt’s reputation, and the bill arrives the day everyone redeems at once. This chapter builds the tether before it builds the claim.

The triad gives us three capacities. To make them tradable without confusing service with money, we need a **unit of account for work**: something that binds energy, hardware, delivery terms, and verification into a transferable typed claim.

Call these **Work Credits (WC)**.

Informally:

A Work Credit is a claim on a standardized unit of triad work (privacy settlement, proof generation, or verified compute) that has been produced and attested under public SLOs.

Key distinction: Work Credits are *energy-priced, not energy-pegged*. They are denominated in work, not in kWh. Energy enters through Facility Energy Receipts (FERs) and VerifyPrice, not through a one-dimensional peg. See [Section 11.7](#) for the full treatment.

Clarification. Throughout this section, “Work Credits” are generic, notional service claims, not product branding and not the base asset. “AI Money” and “proof money” survive only as historical shorthand for service-instrument design spaces; the analytical objects are typed compute, proof, and settlement claims. That separation — claim versus money — is stated once here and defended twice, at [Section 11.4](#) and [Section 11.8](#); it is not repeated as a caveat at every mention.

We can describe Work Credits along a few axes:

11.1 Definition: Work Receipts vs. Work Credits

A common source of confusion is conflating “proof that work was done” with “transferable claim on future capacity.” These are different financial objects. We split them cleanly:

Work Receipt (WR)

A **Work Receipt** is a PIDL artifact proving that a specific unit of work was completed under attested conditions.

- **Content:** Claim hash, proof hash, workload ID, SLA tier, timestamps, hardware profile, prover signature.
- **Properties:** Copyable, verifiable by anyone, *not scarce*.
- **Analogy:** A receipt from a completed transaction. It proves the past but confers no future rights.

Work Receipts are *not* money. They are evidence.

Work Credit (WC)

A **Work Credit** is a transferable instrument issued against Work Receipts under protocol-defined issuance rules.

- **Issuance:** Minted when (a) a valid Work Receipt is accepted by the network, and (b) telemetry confirms SLOs are met.
- **Properties:** Scarce (supply bounded by issuance rules), transferable, fungible within workload class.
- **Rights:** Depend on the design variant (see below).

Instrument	Rights Conveyed	Scarce?	SoV?
Base asset	Bearer fee, settlement, and collateral unit under a constitutional issuance rule. No claim on historical or future work.	Yes	Conditional candidate only
Work Credit / WC-Voucher	Typed claim redeemable or transferable for specified proofs, compute, or settlement capacity at stated workload, tier, place, and delivery terms.	Bounded by DVC	No (service/hedge instrument)

Base asset and Work Credit separation: Hierarchy rule: The base asset is the only object evaluated for monetary premium, and even then only conditionally. Work Credits are useful for service procurement and cost hedging and are explicitly not pitched as stores of value.

Issuance mechanics for Work Credits: For a canonical workload W (e.g., “MatMul of size n with error bound ε ,” “provenance proof for content type C ,” “corridor settlement of size S with anonymity set $\geq A$ ”), define:

- $p(W)$: production cost (energy + hardware amortization + opex) to generate one unit and produce a valid proof.

- $v(W)$: verification cost.
- $r(W) = v(W)/p(W)$: verification asymmetry.

A **Work Credit of type W , tier T** is issued only when:

- A valid Work Receipt for workload W at tier T is accepted by the network.
- **Telemetry confirms** that $\text{VerifyPrice}(W, T)$ and other SLOs (latency, failure rate, decentralization) are within bounds.
- Issuance does not exceed the **issuance envelope** for the current period (see [Section 11.6](#)).

The credit represents only the service rights stated in its terms. A receipt can attest historical work and a voucher can promise future capacity; neither becomes the fee/staking medium or inherits the base asset's conditional monetary candidacy.

11.2 Energy Anchoring

Like PoW, Work Credits are ultimately **energy-anchored**:

- The marginal cost of producing one more credit is bounded below by the energy and hardware required to pass the verification threshold.
- Where SHA-256 PoW has no buyer for its work, this work has one: it powers privacy settlements, proofs of provenance, and AI computation.

This anchoring gives Work Credits:

- **Credible scarcity**: You cannot mint Work Credits without expending real resources to produce proofs and settle flows.
- **Economic meaning**: One credit corresponds to a service somebody was willing to pay for (anonymized payrolls, authentic media, verified inference).

And the cost of the same property. The buyer that gives a Work Credit its economic meaning is also an attack surface the hash puzzle does not have. Demand for SHA-256 work cannot be subsidized away, mandated, or sanctioned, because there is nobody to subsidize, mandate, or sanction; demand for verified inference can be all three. The energy anchor is therefore a service-cost and provenance anchor, not a monetary one. Bitcoin's monetary objectivity and a Work Credit's link to real activity are different properties ([Section 31.1.1](#)).

11.2.1 Layer 0 Maturity and Economic Consequences

Work Credits are only as trustworthy as the hardware that produces them. A world where all proving runs on opaque, vendor-controlled hardware is different from one with diversified, partially open designs.

[Chapter 15](#) (Part III) defines a **Layer 0 maturity ladder** with grades L0-A through L0-D. Here we preview how those grades affect Work Credit economics:

L0 Grade	Hardware Trust Level	Economic Treatment
L0-A	Closed hardware + attestation + multi-party audits + diversity	WC accepted at full value if diversity thresholds met; risk premium priced into fee schedules
L0-B	Lot sampling + imaging + bounded side-channel budgets	WC accepted at full value; lower risk premium
L0-C	Partial open (open RTL for critical components)	WC may qualify for “open-profile” premium or priority tiers
L0-D	Fully open designs	Highest trust tier; may command premium in markets that value sovereignty

How this affects issuance and pricing:

1. **Tiered issuance caps:** WC minted on L0-A profiles may face stricter issuance limits than L0-C/D profiles. This prevents the system from becoming dependent on opaque hardware.
2. **Market pricing:** WC from different L0 grades can trade at different prices if markets distinguish them. Treasuries and institutions may pay premiums for L0-C/D-backed capacity.
3. **Risk disclosure:** Every Work Receipt includes the hardware profile (HID) used. Aggregated telemetry shows what fraction of total WC is backed by each grade.
4. **Deprecation impact:** If an L0-A profile is compromised (TEE backdoor discovered), WC minted on that profile can be quarantined, discounted, or excluded from certain uses. This is the “bond downgrade” analog.

Why this matters for service-claim quality: If 90% of Work Credits are minted on L0-A hardware and a major TEE is compromised, claim quality and DVC take a hit—not because the cryptography failed, but because the service path depended on a common hardware assumption.

By making L0 grade explicit and tying it to economic consequences, the system:

- Creates incentives to invest in open hardware.
- Bounds the impact of hardware compromises.
- Gives users and allocators the information to price risk appropriately.

This is covered in detail in [Section 15.4](#) (Part III). The key point for Part II: **hardware trust is not binary; it is graded, measured, and priced** — and because it is priced, the service claim’s quality is legible without anyone having to take the word “typed” on faith.

11.3 Robots, AI, and the Demand for Work Money

In a robot- and AI-heavy economy the demand story for Work Credits becomes almost embarrassingly straightforward.

Consider a warehouse or factory where most of the physical activity is carried out by robots, and most of the planning and oversight is handled by models. The day-to-day budget splits into:

- Energy to power robots and data centers.
- Compute to run models and generate proofs.
- Privacy and settlement to pay workers, suppliers, tax authorities, and investors without leaking trade secrets or exposing everyone's graph.

Each of those budgets expresses itself as a recurring need for receipts:

- Proof of correct inference for high-stakes decisions.
- Proof of compliance and risk calculations for regulators and insurers.
- Privately settled wages and vendor payments with lawful audit trails.

Today that flow is mediated through cloud bills, payroll files, bank wires, and audit PDFs. In the stack described here, it can be mediated through Work Credits and receipts. The robots and services are paid in typed claims; in exchange they produce receipts that can be verified cheaply and settled over privacy rails. High turnover can make the claims useful working balances. It does not establish reserve demand for them or for the base asset.

This is one concrete way to discipline remarks like “money will just be energy.” Energy, compute, and hardware are inputs; Work Credits denominate future access to verifiable, robot-mediated output. They do not become joules, reserve assets, or monetary anchors by denomination.

11.4 Why Work Credits Remain Service Claims

A claim can be honest and still not be money. Scarcity, transferability, public verification, censorship resistance, and recurring demand are together enough to make a Work Credit a high-quality commodity or forward service claim; none of them, singly or jointly, is enough to make it savings. The reason is physical rather than rhetorical: the credit stays exposed to workload basis, location, hardware, SLA, expiry, delivery, and common-cause risk, and a promise of future service may carry duration. It is a warehouse receipt with a date on it, and a savings instrument has to survive the day the warehouse is the problem. DVC and the evidence artifacts determine whether the claim is fulfillable; the Native Monetary Buyer Map determines whether the separate base asset has a holder constituency capable of bearing loss. Two different questions, two different instruments, and neither answer transfers.

11.5 Monetary Role

Work Credits sit at the junction between **base capacity** and **service contracting**:

- For operators, they are **revenue in kind**: miners/provers/routers earn credits by contribut-

ing triad capacity.

- For users, they can be **pre-paid capacity**: hold credits to secure future access to specified triad services, or trade expected scarcity.

Their service value depends on:

- **Demand for specific workloads**: Work Credits tied to high-value workloads (e.g., compliance proofs, LLM inference) may command higher premia.
- **Governance & telemetry honesty**: If VerifyPrice and decentralization metrics are falsified or gamed, the link between credits and real work weakens.
- **Regime pressure and deliverability**: Differential demand may rise as substitutes weaken, while DVC may fall; the Pressure–Capacity Corridor prevents a monotonic crisis claim.

In portfolio terms, Work Credits are typed service or capacity exposure. LP and staking positions are the equity-like or derivative layer; the base asset is evaluated separately.

11.6 Issuance: Tying Credits to Real Capacity

There are many possible issuance schemes; what matters here is not choosing a particular curve, but enforcing two principles:

1. **Issuance is legible.**
2. **Issuance is constrained by real capacity.**

One extreme is the Bitcoin model: a fixed schedule, regardless of demand, with the understanding that price will equilibrate. Another extreme is a pure capacity-linked model: Work Credits are minted only when new proving, compute, and settlement capacity comes online and is registered with telemetry; they behave almost like tokenized capacity reservations. In practice, a hybrid is likely: a predefined issuance envelope over time, modulated by capacity growth and burn.

In such a system, adding a new proving cluster, plant, or corridor is not just a marketing slide; it is an event that expands the envelope of Work Credits the system can credibly support. The converse is also true: if plant retires or corridors die and are not replaced, issuance that continues on autopilot will show up as VerifyPrice drift, SLA breaches, and deteriorating energy metrics. The governance layer’s job is not to guarantee any particular price, but to keep issuance and capacity in rough proportion and to make any departures visible.

From an allocator’s perspective this yields a familiar pattern: Work Credits look like equity in (or claims on) a portfolio of infrastructure (proving farms, AI chains, privacy corridors, and plants) rather than a purely arbitrary balance sheet. The difference from conventional “infra tokens” is the insistence on receipts and KPIs: if claimed capacity and observable behavior diverge, the discrepancy is not a rumor; it is a datapoint.

11.6.1 The Supply Balance Equation

Here is the question that kills most “energy-backed” tokens in the cradle: **why does demand not simply expand supply, neutralizing scarcity?** The mechanism sounds self-defeating. If Work Credits are minted whenever verified work is done, then every new customer who arrives mints new supply to serve them. The token is busy, the network grows, and the scarcity — the very thing that was supposed to make holding worthwhile — evaporates into its own success. Demand, on this reading, is not a source of scarcity but an exhaust valve for it.

That is not a bug in the design. It is the design, correctly read. Work Credits are minted against delivered service, so supply growing with demand is not leakage; it is the claim staying honest. A service claim whose supply could not grow to meet service demand would be a shortage instrument, not a savings instrument — and neither of those is a reserve asset. The interesting quantity is therefore not gross issuance but what remains after the service is delivered, and that is a balance, not a schedule.

The answer lies in the **net supply dynamics**. We can express this as a balance equation:

$$\Delta \text{Outstanding WC} = \text{Issuance} - \text{Redemption/Retirement} \pm \text{Governance}$$

Lost or forgotten wallets do not leave the outstanding count — the units still exist on-chain — so they are not a term in this identity; the *effective* (circulating) supply additionally nets them out.

Term	Definition	Typical Magnitude
Issuance	New WC minted against Work Receipts, subject to issuance envelope	Bounded by schedule or capacity ceiling
Redemption/retirement	WC destroyed when the specified service is delivered or the claim expires	Set by contract terms
Lost	WC in lost/forgotten wallets — retained in <i>outstanding</i> , excluded from <i>effective</i> supply	~1–2% of supply per year (empirical from BTC)
Governance	Adjustments via protocol upgrades (rare, requires supermajority)	Near zero in steady state

For Work Credit supply to remain aligned with deliverable service, the design must ensure:

$$\text{Outstanding claims} \leq \text{haircut stress-adjusted DVC over the delivery horizon.}$$

Read that as a solvency test, not a scarcity test. It is the warehouse being asked to show

that every receipt outstanding corresponds to something physically present and deliverable. It places no ceiling on how many receipts a busy system writes; it places a ceiling on how many can be outstanding at once, and the ceiling moves only when the physical world moves. That distinction carries the whole chapter:

- **During growth:** Work Credit issuance may grow with independently verified DVC and contracted demand.
- **At maturity:** Claims retire through delivery, expiry, or cancellation under published terms. Scarcity can affect service price without making the claim a reserve asset.

What prevents runaway issuance? The same four doors, in the order an operator would try them:

1. **DVC envelope:** Total WC outstanding for each workload, tier, and horizon is capped by haircut stress-adjusted DVC.
2. **Capacity evidence:** The envelope expands only when FERs, FCRs, hardware profiles, and scenario flow support additional deliverable service.
3. **Contract retirement:** Delivery, expiry, cancellation, and default rules remove or impair claims explicitly.
4. **Governance friction:** Changes to service-claim rules require supermajority and are visible in dashboards before activation.

Telemetry that detects supply risk:

- **Inflation rate:** Issuance / circulating supply. Should decline over time.
- **Retirement rate:** retirements / issued. Should track delivered service and expiry terms.
- **Net supply change:** Δ Supply per epoch. Positive during growth, flat or negative at maturity.
- **Issuance vs. capacity:** If issuance grows faster than verified capacity, this signals potential dilution.

If these metrics drift outside healthy ranges, it becomes visible in dashboards—holders and operators can respond before the SoV thesis is undermined. A system that mints against a warehouse cannot hide a shrinking warehouse; the drift shows up in the issuance numbers before anyone has to allege anything.

11.7 Energy-Priced, Not Energy-Pegged

In popular discourse one often hears that “money will just be energy,” especially in the context of robotics and AI. As shorthand, this is attractive: robots and data centers run on electricity, so why not quote everything in kWh and be done with it? The problem is that not all kilowatt-hours are created equal. Time of day, grid node, reliability, carbon intensity, and

siting constraints all affect their economic and political meaning. A winter-peaking kWh on a stressed urban grid is not the same as a curtailed hydro kWh in a remote valley. If we pretend otherwise, we smuggle a lot of hidden politics and risk into the unit of account.

The discipline in this thesis is to admit that complexity and route around it with **receipts**.

Energy is measured and wrapped into Facility Energy Receipts (FERs), which record, for each facility and time window:

- kWh in,
- kWh delivered to IT,
- heat reused,
- PUE/ERE/WUE,
- water use,
- carbon intensity,
- and outages/curtailments.

On top of that, we measure **verified work per kWh**:

- η_{VFLOP} for FLOPs,
- proofs-per-kWh for proof workloads,
- swaps-per-kWh for settlement.

VerifyPrice then tells us how much it costs, in time and money, for an independent verifier to check that a given amount of work was done.

The result is an implicit conversion path:

energy → FERs → **verified work (proofs/FLOPs/swaps)** → receipts → **Work Credits**

We never pretend that *one Work Credit is one kilowatt-hour*; instead we make it easy for anyone to estimate, at any given time, **how many kilowatt-hours of which quality and where in the world** sit behind a portfolio of Work Credits, via the receipts.

Pricing in energy then becomes an **inference problem for markets**:

- Work Credits are implicitly energy-priced because their production and redemption depend on energy-intensive workloads whose cost curves are public.
- FERs and VerifyPrice make those curves legible without forcing a brittle kWh peg.

What distinguishes Work Credits from naive “energy tokens” is precisely this separation:

- The unit of account is denominated in **work**, not in kWh.
- Energy enters through FERs, η_{VFLOP} , and VerifyPrice, not through a one-dimensional peg.

Markets, regulators, and builders can look at those receipts and say, with some confidence, “a Work Credit currently corresponds to about this much capacity, with this energy and carbon profile, under these SLAs.” That is enough to make the asset priceable and analyzable without forcing it into a crude energy standard.

11.8 Why Work Credits Are Typed Service Claims, Not Presumed Money

It is equally easy to make the opposite error: to observe that Work Credits are scarce, transferable, and tied to necessary workloads, then promote them from service claims into money by vocabulary. This thesis no longer makes that move. A Work Credit is a typed claim on capacity or service whose quality depends on workload, location, hardware, SLA, delivery window, and the stressed path that can actually fulfill it. That case was made at [Section 11.4](#); this section asks the harder question it left open. What would have to be *true* — not said, not branded, but observable — for a Work Credit to cross from service claim into money?

Four things, each visible in telemetry before it is visible in price:

1. **Issuance is constrained by stress-deliverable capacity, not nameplate capacity.**

FERs, FCRs, and Delivered Verified Capacity bound claims by the service that survives scenario-specific minimum cuts. Gross installed hardware, energy procurement, or benign-state output is insufficient.

2. **Demand can be budgeted without being monetary.**

Proofs, inference, settlement, and related workloads may sit on recurring OPEX lines. That supports a service market. It does not establish reserve demand, long holding periods, self-custody, loss-bearing capacity, or countercyclical accumulation.

3. **Verification makes the claim legible, not monetary.**

A PIDL receipt can make fulfillment cheap to check and FCR/FER artifacts can make provenance auditable. Evidence reduces information asymmetry; it does not erase delivery, basis, counterparty, or temporal risk.

4. **The claim must not absorb project duration by implication.**

A prepaid capacity claim, a claim on future service, and a project note are different instruments. If a Work Credit promises par, redemption, a coupon, or emergency support against construction failure, it has become duration-bearing credit and must be labeled and underwritten as such.

So what would the crossing actually look like? Not a marketing decision and not a milestone on a roadmap, but a change in *holding behavior* that the dashboards can see: holding periods lengthening rather than claims being redeemed at term, balances retained through drawdowns instead of recycled into service, and — the tell that would settle it — buyers acquiring the asset *while the service it claims is cheap and available*, which is the signature of savings rather than procurement. None of that is a design property. It is holder behavior, and it is exactly what the Native Monetary Buyer Map watches for on the base asset ([Section 24.8.8](#)). Until those series move, Work Credits belong beside capacity vouchers and service contracts, not beside the base monetary candidate. They may trade at a premium for scarce delivery rights and may serve as collateral where their terms permit. Those are properties of typed

service claims. Monetary treatment remains a separate empirical question answered, if at all, by the holder-side mechanism and buyer map (Section 11.10.8).

11.9 Economic Linkage: How Triad Demand Becomes Asset Value

This subsection answers the question that separates a *systems manifesto* from a *monetary thesis*: **Why does demand for Privacy, Proofs, and Compute raise the value of *holding* the asset, rather than merely rewarding *consuming* a service?**

The answer has four parts.

1. What exactly is the asset? The **asset** in this framework can take several forms, but all share a common structure:

- **Network tokens:** Native units of the protocol (analogous to ETH or BTC) that are required to pay fees, post collateral, and participate in governance.
- **Work Credits:** Claims on verified capacity, minted against energy-anchored work.
- **Corridor/Pool shares:** LP positions or staking rights in specific privacy corridors, proof factories, or compute networks.

What matters is that **all uses of the triad must flow through the asset**. You cannot get a proof, settle a private payment, or buy verified compute without either holding or acquiring the native unit.

2. Why is it scarce in a monetary sense? Scarce *capacity* does not automatically imply a scarce *asset*. The link is forged through:

Mechanism	How It Creates Monetary Scarcity
Capped issuance	Total supply follows a predefined schedule (e.g., halvings) or is bounded by capacity growth, not governance fiat.
Fee burns	A portion of every fee is permanently destroyed, removing units from circulation as usage grows.
Collateral lockups	Provers, routers, and LPs must bond assets to participate; this removes circulating supply proportionally to network activity.
Energy-anchored minting	New credits are issued only against verified work backed by FERs; you cannot mint by decree.

The result: **issuance is bounded by physics (energy, hardware), while usage acts differently on the two instruments — the base asset shrinks with usage through fee burns, and Work Credits retire as service is redeemed or claims expire**. Demand is driven by compulsory workloads. This is the scarcity structure of a commodity, not an IOU.

3. Why does demand for triad capacity raise the value of holding the asset? The naive chain runs: demand → users must acquire tokens to pay fees → fees paid, a portion burned → burns reduce supply while demand rises → “increased demand + reduced supply → price appreciation” → holders capture the triad’s economics. Steps 1–4 are routing mechanics and survive review. Steps 5–6 are the tokenomics inference, and [Section 11.10.7](#) withdraws them: a burn is a buyback, a fee stream is a DCF input, and neither confers money. The sentence this section will stand behind is narrower: the holder is buying a competitively priced claim on a fee stream bounded by Δ , plus a collateral floor. Where any premium beyond that could come from is deferred to [Section 11.10.8](#).

More concretely:

- **Fee revenue:** Every proof, every private settlement, every verified inference pays a fee denominated in the native asset. This routes service demand through the asset — the routing claim only; [Section 11.10](#) bounds what the fee can be.
- **Burn mechanics:** A fraction (30–50% in the reference design) of fees is burned. Whether burns exceed issuance at steady state is an empirical outcome the fee-coverage boards measure, not a design guarantee.
- **Required collateral:** Provers, routers, and LPs must stake tokens proportional to their capacity. This locks supply and aligns incentives with network health; [Section 11.10.9](#) bounds how much lockup the system can carry.
- **Priority/governance rights:** Holding grants access to premium SLA tiers, governance votes, and first-mover allocation of scarce capacity.

Key insight: The holder is not buying “exposure to price” (reflexive speculation). The holder is buying **a share of the fee stream from indispensable workloads**, denominated in an asset whose *issuance* is bounded by physics. Whether burns exceed issuance at steady state — whether net base-asset supply falls as those workloads grow — is an empirical outcome the fee-coverage and issuance boards report, not a design guarantee. This is closer to equity in a utility than to a collectible.

4. What prevents capacity providers from capturing all value while holders get diluted? This is the classic “utility token trap”: if operators earn all the fees and governance can inflate supply, holders are just exit liquidity.

The stack avoids this through:

The economic design goal: routing integrity — fees flow to retirement, staking, and operator share in fixed proportions, and issuance beyond verified capacity is visible in telemetry. Whether that routing produces flat or declining net supply is an outcome the fee-coverage and issuance boards report, not a promise made here.

Summary: the Value Capture Lemma in plain language

Risk	Mitigation
Operator rent extraction	Fees are split: burn + stakers + provers. Operators cannot capture 100%; a structural portion goes to asset retirement.
Governance inflation	Base-asset issuance follows the constitutional schedule, independent of Work Credit dynamics; Work Credits are minted only against verified capacity (FERs, telemetry), and minting beyond real capacity triggers SLO breaches visible in dashboards.
Holder dilution	Base-asset issuance is constitutionally scheduled; burns offset it as an empirical outcome the boards report, not a design promise.
Value leakage to fiat	Core operations (fees, collateral, rewards) are denominated in the native asset, not fiat. Fiat is an off-ramp, not the unit of account.

Demand for triad capacity (Privacy, Proofs, Compute) becomes store-of-value premium for the native asset **only if** five conditions hold at once: all usage requires the asset, usage retires supply, capacity provision requires staking, issuance is constrained, and capacity cannot be bypassed. The formal statement, with each condition specified, is the [Value Capture Lemma](#) that follows; it is not re-listed here.

These conditions are necessary, not sufficient. If any one of them fails, the system may be indispensable infrastructure and still not a store-of-value asset. Nothing here establishes the converse, and the thesis does not claim it.

A related but separate claim. The five conditions are properties of the monetary design; they say nothing about the durability of the demand the design is routing. That is a distinct empirical claim—that workload demand is *structural* (AI, commerce, and compliance budgets rather than hype cycles)—and the thesis treats its failure as its own falsifier rather than as a condition of the lemma: a workload mix that becomes and remains predominantly speculative retires the store-of-value claim through Red Line 5 ([Section 28.6](#)), not through this lemma.

Lemma 11.1 (Value Capture Lemma). *Demand for Privacy, Proofs, and Compute creates store-of-value premium for the native asset **only if** five conditions hold simultaneously:*

1. **Required fee medium:** *the asset is required for core fees, and users cannot pay in fiat, stablecoins, or other tokens at equivalent service quality.*
2. **Supply reduction:** *a meaningful share of fees is burned or permanently retired.*
3. **Collateral lockup:** *operators must stake the asset as collateral to provide services.*
4. **Issuance discipline:** *issuance is capped by schedule, capacity, or both—not by governance fiat.*
5. **Non-bypassability:** *users cannot obtain equivalent triad capacity through bypass channels (cloud contracts, stablecoin-denominated services, direct fiat payment to operators) that avoid touching the asset.*

*If any condition fails, the system may be useful infrastructure but no **monetary premium** accrues — the asset may still hold ordinary discounted-cash-flow value. This lemma is the bridge between “these capacities are indispensable” and “therefore a specific asset accrues the value of that indispensability.” It is a necessary condition for the monetary claim and not the monetary claim itself (Section 11.10.7, Section 11.10.8).*

One reconciliation with Section 11.10.8 belongs here. The five conditions are jointly sufficient for premium and individually necessary; within that set, **issuance discipline** is the one whose failure the others cannot compensate for. And Section 11.10.8 is not an alternative gate: it describes *where* the premium comes from once this gate is open — a holder-side service flow — not a second, competing path to moneyiness that routes around these conditions.

This is the bridge between “these capacities are indispensable” and “therefore a specific asset, rather than its operators or its customers, captures the value of supplying them.”

What the five conditions do not address. Every one of them governs *where* fees go. None governs *how large* they can be, and a routing rule applied to a fee of nothing captures nothing. Section 11.10 takes up the level question, finds that the standard objection to taxing a contestable market is wrong for a reason worth stating, and then finds that the correct answer supports a narrower conclusion than this Part has been drawing. The conditions themselves are unchanged.

The lemma has a frequently-missed converse. It states the conditions under which native demand accrues *economically* to the asset. It does not state that movements in the asset’s *observed market price* are caused by that native demand. Those are different machines, and conflating them is the most common way a monetary thesis fools its own author.

Corollary 11.1 (Market-Price Non-Equivalence). *External financial wrappers can create price demand without exercising privacy settlement, purchasing proofs, consuming verified compute, paying native fees, locking collateral, or producing burns. Conversely, wrapper redemptions, leverage unwinds, dealer hedging, or broad risk-off flows can depress the asset’s price while native monetary usage improves. Therefore*

$$\text{price appreciation} \not\Rightarrow \text{monetary adoption}$$

and

$$\text{monetary adoption} \not\Rightarrow \text{immediate price appreciation.}$$

*Price is an output to explain, not a protocol KPI to worship. The monetary thesis must be evaluated through native fees, burns, collateral, settlement, receipt volume, verification health, reachability, and non-custodial use. Market-price interpretation requires a separate financial-flow telemetry layer (**VerifyFlow**, Section 24.8).*

11.9.1 Value Capture vs. Price Capture

It is worth separating the two phenomena explicitly, because the thesis is routinely “confirmed” by the wrong one.

Value capture is the existing loop: triad usage → native fees → burns, collateral lockup, and operator demand → scarcity and *cash-flow accrual* to the asset. It is governed by [Lemma 11.1](#). What it produces is a claim on a fee stream plus a collateral floor; whether the asset also carries monetary premium is a separate question, answered separately in [Section 11.10.8](#).

Price capture is demand for exposure to the asset’s price, arriving through spot ETFs, exchange custody, treasury companies, margin products, options, futures, swaps, leveraged and inverse ETPs, passive indices, and systematic trading rules. Price capture can occur with no protocol use whatsoever.

Native value capture	Price capture	Interpretation
Strong	Strong	Genuine adoption plus favorable market realization
Strong	Weak/negative	Protocol improving while wrappers or risk markets sell
Weak	Strong	Financialized speculation or wrapper-led adoption
Weak	Weak	Failed or immature monetary thesis

Table 11.1: Native value capture and observed price appreciation are neither equivalent nor mutually necessary. Only the top-left cell is monetary validation.

The diagonal is what makes the thesis falsifiable. A thesis that treats every price rise as confirmation and every price fall as noise cannot be wrong, and therefore cannot be right either.

11.10 Fee Incidence and the Level of the Fee

Every condition of the [Value Capture Lemma](#) answers one question: *where does the fee go?* Required medium, supply reduction, collateral lockup, issuance discipline, and non-bypassability are all routing rules. None of them answers the question a valuation actually needs answered, which is *how large the fee can be*. A protocol that routes the whole of a fee of nothing into a burn address has satisfied five conditions and captured zero.

The gap has a sharp form, and it reached us as a fatal objection rather than as a note about an omission:

“You have deliberately engineered the most contestable market you could. Workloads are canonical and interchangeable, verification is cheap enough that no trust is required, admission is open, and

there are no vendor chokepoints. A contestable market competes economic profit to zero. Zero profit means no surplus, and no surplus means nothing to burn."

The objection is wrong. It is worth showing exactly why, because the correct answer rescues considerably less than it first appears to, and what survives is not the claim this Part started with.

11.10.1 Incidence: A Fee Is a Wedge on Turnover, Not a Claim on Profit

The objection conflates two different bases. A protocol fee is levied on *transactions*, not on *residual profit*. It is a tax wedge, and tax wedges do not require anybody to be earning rents.

Take the market the thesis is trying to build: many providers, free entry, roughly constant returns to scale in the relevant range, a canonical workload so that units are fungible. Long-run supply is then approximately horizontal at marginal cost c (equivalently the protocol's all-in unit cost c_p of Section 11.10.3; write $c \equiv c_p$). Impose a per-unit protocol fee f on that market and the supply curve the buyer faces shifts to $c + f$. The new equilibrium has buyers paying $c + f$, providers receiving c , and providers earning exactly the zero economic profit they earned before. The protocol collects $f \cdot Q$. Nothing about that revenue came out of operator margin, because there was no operator margin to take it from; in the long-run competitive case the entire statutory burden lands on the buyer.

Operator margin compression is irrelevant to fee revenue. The fee is a wedge on turnover. Turnover is what a contestable market maximizes.

This is standard incidence analysis, and it means the objection proves too much: if contestability destroyed fee bases, no commoditized market could ever be taxed, and every one of them is.

Incidence explains who *bears* a fee. It does not explain why a positive fee is *sustainable*, and the two should not be run together. Exchanges, clearing houses, and toll roads are lucrative on top of viciously competitive flows for a further reason: the flow is competitive but the fee layer is not. There is one place to clear the contract, one road across the water, one book where the trade is legally final — so the competition among users of the layer never becomes competition to *be* the layer. That inescapability is a property those venues possess and a protocol must earn; it is the subject of Section 11.10.3 and Section 11.10.6, and the fork channel is precisely the case where it is not earned.

The short run is different, and the difference matters. With installed capacity fixed, supply is not horizontal. Deployed hardware has already sunk its capital, so in the short run the supply curve slopes upward and part of the fee *is* absorbed as a reduction in quasi-rent to

existing machines rather than passed to buyers. That is not a problem for this quarter's fee revenue; it is a problem for the next round of capacity. Quasi-rent is what funds replacement, and an operator whose realized return on deployed hardware is being clipped by a fee re-prices the entry decision accordingly. The long-run pass-through result therefore describes the destination, not the transition, and the transition is where entry and capacity growth are decided. A fee schedule set as though incidence were instantaneously long-run will suppress the capacity it is trying to tax.

11.10.2 The Base, and Why It Deflates

Fee revenue is not f . It is

$$\text{Rev} = f \cdot Q(c_p + f),$$

and both terms are endogenous. Two consequences follow, and the second is the more serious one.

Quantity responds to the wedge. The tollbooth analogy holds only where there is no parallel free road. Raising f raises the delivered price, and at the margin some volume declines, defers, batches, or leaves. What the protocol can extract is bounded not by what buyers would pay for the service but by what they would pay *through this channel rather than the next-best one*, which is the subject of [Section 11.10.3](#).

An ad valorem fee on a deflating unit price is a shrinking real toll. For this paragraph the protocol's fee is modeled as *ad valorem* — a share of unit price rather than the per-unit wedge f used in the incidence analysis above; under a per-unit fee only the volume term below moves, and the deflation mechanism this paragraph isolates does not arise. [Section 10.1](#) argues that verified compute stays *scarce* even as raw compute deflates. Grant that argument entirely; it does not touch this one. Even if the verification premium is fully durable, a fee expressed as a percentage of a unit price that falls secularly collects fewer real resources per unit of service delivered every year. Real fee revenue then grows only if

$$\underbrace{\text{growth in verified units delivered}}_{\text{volume}} > \underbrace{\text{rate of decline in real unit price}}_{\text{deflation}}.$$

That inequality is an empirical bet, not a theorem, and the thesis is making it. It may well be a good bet — the historical pattern in compute is that unit-cost decline is met by more than proportional volume expansion — but the thesis has nowhere stated it as a bet, measured it, or said what would settle it against us. Condition B of Red Line 14 ([Section 28.6](#)) exists to settle it: deflation-adjusted fee-plus-burn turnover that fails to grow while physical throughput grows is this bet losing, in public.

11.10.3 The Bound on the Sustainable Fee

Incidence tells us the fee has a base. It does not tell us how high the fee can be set. That ceiling is not a function of operator margin either; it is a function of what the protocol channel is *differentially worth*.

Differential Value Δ and the Sustainable Fee Bound

Let c_p be the protocol's all-in cost of delivering canonical workload W at tier T , and c_b the all-in cost of the best bypass channel delivering the same computational content without protocol-grade verification. Let Δ be the marginal buyer's willingness to pay for the protocol's differential properties: cryptographic checkability, neutrality, credible non-discretion, and non-custodial exit.

The buyer routes through the protocol only if $c_p + f - \Delta \leq c_b$, so the sustainable fee satisfies

$$f \leq \Delta - (c_p - c_b)$$

Here Δ is a scalar wedge in price units, not the difference operator used elsewhere in this Part.

The first term is the one a protocol designer wants to think about. The second is the one that decides the outcome, and it is **structurally positive**: verified compute costs *more* to produce than unverified compute, because proving overhead and redundancy are real resources spent on top of the underlying work. So the protocol begins every negotiation with a cost handicap and must recover it out of Δ before it collects a single unit of fee.

The document's own tier multipliers are cost pass-through, not evidence of willingness to pay. Section 20.3 sets fee multipliers of $1.5\times$ for Silver and $2.5\times$ for Gold. Read the same table's redundancy column: Silver is $2\times$ verification, Gold is $3\times$ verification plus audit. The multipliers are rationale-set against that redundancy cost structure — intended to cover redundancy plus margin — but the document never establishes that they do: $1.5\times$ against $2\times$, and $2.5\times$ against $3\times$ -plus-audit, leave the coverage question open by construction, and it is a question about c_p and the margin above it, not about demand. They are a statement about c_p , and nothing in the document establishes that any buyer's Δ at Gold tier is $1.5\times$ larger than at Silver. If it is not, then headroom *narrows* as tiers rise, and at high tiers $\Delta - (c_p - c_b)$ can be negative — a tier that is technically superior, honestly priced, and unsellable. That is a testable prediction against the thesis and it is now measured (Section 28.6, Red Line 14), because it should not have been left as an inference from a table caption.

11.10.4 Bypass Is a Magnitude, Not a Binary

Condition 5 of the [Value Capture Lemma](#) is written as an existence condition — users *cannot* obtain equivalent capacity through a bypass channel — and Red Line 6 is written as a direction condition: usage grows while native fees do not. Neither is wrong, and neither is changed here. But both are read, in practice, as though bypass were a switch.

It is not. Bypass is never impossible. It is priced. There is always a channel that delivers the same computational content without protocol-grade verification, and the only question is what the buyer gives up by using it — which is exactly Δ . Condition 5 should therefore be *read and measured* as a magnitude claim: not “no bypass exists,” which is false everywhere, but “ Δ is large enough, for enough of the volume, to support a fee worth burning.” The condition stands as stated. Its *interpretation* is what this section tightens, and its measurement is what Red Line 14 supplies.

This is why the chain-strength rating for non-bypassability is reduced from Strong to Medium in this version ([Table 7.1](#)). The condition is well specified as a binary and the binary is well enforced. The economics requires a magnitude, and the magnitude is nowhere estimated. A rating of Strong was a rating of the enforcement machinery, not of the quantity the argument depends on.

11.10.5 What Determines Δ , and Why It Is Probably Thin

To first order, the differential value of a proof is the loss it prevents, times the probability that loss survives the alternative:

$$\Delta \approx \varrho \mathcal{L},$$

where $\mathcal{L}$ is the loss from undetected misexecution or denial of service and ϱ is the residual probability of that loss under the best bypass channel. This operationalization *lower-bounds* Δ : it prices only loss-prevention, while the full definition in [Section 11.10.3](#) covers four properties — cryptographic checkability, neutrality, credible non-discretion, and non-custodial exit — so neutrality and non-discretion enter as additional terms the $\varrho \mathcal{L}$ form omits. (ϱ and $\mathcal{L}$ are used rather than the conventional π and L , which denote the inflation rate in [Section 3.1](#) and the leverage multiple throughout [VerifyFlow](#) and [Appendix H](#).)

Δ is a distribution, not a number. Before evaluating either term it is worth fixing what object Δ is. Every buyer has their own ϱ and their own $\mathcal{L}$, so Δ is a distribution across the buyer population, and the fee bound of [Section 11.10.3](#) binds at the *marginal* buyer — the last one the protocol wants to retain at the posted fee — not at the median one. These come apart, and the difference is the whole commercial question. A thin median is entirely consistent with a lucrative fee, provided the upper tail is deep enough to be worth serving: counterparties nobody will indemnify because they are sanctioned or pseudonymous, flows crossing juris-

dictions where no single forum will hear the claim, regulated model-risk and audit functions that must evidence correctness to a supervisor rather than merely believe it. For those buyers ϱ is not small, because the indemnitor substitute is unavailable rather than merely inferior. The protocol's commercial problem is therefore not "is the median workload willing to pay," which it plainly is not, but "is the high- Δ tail large enough, and reachable enough, to carry a fee." That is a sizing question the thesis has not answered, and Red Line 14 measures it capacity-weighted across the served mix rather than assuming it.

And at the median the answer is unflattering. Written as $\varrho \mathcal{L}$, the critical term is ϱ , and here the thesis has been flattering itself. The bypass channel is not "unverified." It is *reputationally and contractually* verified. A hyperscaler's alternative to a cryptographic proof is its balance sheet, its SOC 2 and ISO audit reports, its published incident history, and an indemnity clause enforceable in a court that will actually hear the case.

A creditworthy, suable indemnitor is an excellent substitute for a cryptographic proof.

For an ordinary commercial workload in a benign state, the buyer's exposure to undetected misexecution is not $\mathcal{L}$; it is $\mathcal{L}$ net of a contractual recovery from a solvent counterparty under a legal system that functions. That makes ϱ small, and therefore Δ thin over the bulk of the distribution. Thin Δ against a structurally positive $(c_p - c_b)$ is the strongest form of the objection this section opened with, and it is not answered by incidence analysis. It has to be stated plainly, because a critic will otherwise state it first and more crudely: **for the median commercial workload in a working legal order, cryptographic verification is a premium product competing against a cheaper substitute that is good enough.** What that sentence does *not* say — and what a fair reading of it must preserve — is that the median is the marginal buyer. It is not, and the fee is set against the tail.

11.10.6 Two Bypass Channels the Thesis Has Not Been Counting

Both of the following erode the base without tripping the instruments currently pointed at bypass.

Partial bypass at the intensive margin. Bypass has been modelled as a routing decision — the buyer uses the protocol or does not. In practice the buyer uses *less protocol per unit of underlying activity*: batching many operations under one proof, aggregating proofs recursively, or using the protocol purely for attestation while executing the work elsewhere. Each is a legitimate engineering optimization, and each cuts native fees per unit of real economic activity. None of it registers as a Red Line 6 style "usage without fees" signal, because usage does still route through the protocol — just thinner. The correct denominator for fee intensity is the un-

derlying economic activity being attested, not the count of protocol interactions, and the thesis has not been measuring it that way.

The protocol can be forked. This is the more important omission. In an open-source, open-admission, canonically-specified market, the cheapest bypass channel is not a hyperscaler. It is *this protocol, copied, with a lower fee and no burn*. Every property the thesis is proud of — published workload specifications, open verifier implementations, permissionless participation, no vendor chokepoints — lowers the cost of producing that copy. A fork inherits the technology and discards the tax.

What actually prevents the fork from winning is not any of the five conditions. It is liquidity depth in the native asset, anonymity-set size in the privacy corridors, the installed collateral base and the operator relationships around it, the Schelling-point status of the workload registry, and the accumulated receipt history that makes one chain's attestations worth citing. Those are network effects and coordination assets. They are not enforcement rules, and the thesis has been attributing the moat to enforcement rules. Stated honestly: **Section 11.12 explains why a user cannot get the service without the asset on this protocol. It does not explain why the user cannot get it on the copy.** Whatever answer exists is a liquidity and coordination argument, it is empirical, and it belongs to link 7; whether holders then warehouse the asset's risk belongs to link 8.

This is not purely hypothetical, and the one well-documented case ran exactly this experiment against exactly this kind of moat. In 2020, SushiSwap copied Uniswap's open contracts and paid liquidity providers to migrate — a deliberate attack on a liquidity-and-coordination moat by an identical copy with better short-term incentives. The moat held. Liquidity depth, accumulated positions, and the receipt history a trader's history lives in did not transfer with the code, and the copy won a share rather than the market. That is one observation in one corner of the design space, not proof — but it is evidence that the moat in question is real, and evidence about *which* moat: what defended the incumbent was not its technology, which was copied freely, but the coordination assets stacked on top of it. It is precisely those assets, and not the open specification, that a fork of this protocol would have to defeat.

11.10.7 What the Five Conditions Establish, and What They Do Not

Collecting the above, the lemma survives with its scope corrected downward.

What they establish. That a contestable market can be taxed and the proceeds routed to a token. Specifically: (i) a competitively priced claim on a cash-flow stream $f \cdot Q$, bounded above by the per-unit wedge Δ applied to volume — $(\Delta - (c_p - c_b)) \cdot Q$ at the sustainable-fee bound of Section 11.10.3 — and deflating with unit price; and (ii) a balance-sheet floor from collateral required in the native unit. Both are real. Both are also, precisely, quantities that a

discounted-cash-flow valuation captures — a fee stream and a lockup.

What they do not establish. Monetary premium. A monetary premium is by construction value *in excess of* discounted cash flows; it is what gold has and a pipeline does not. Section 11.9 states the holder as buying “a share of the fee stream from indispensable workloads,” “closer to equity in a utility than to a collectible.” That description is accurate — and it is why this section withdraws the tokenomics inference made above it. A utility equity earns its cost of capital; it does not carry monetary premium. The two registers now stated side by side in this Part — routing rules on the one hand, price-mechanics claims on the other — are separated deliberately: the routing rules are design commitments that survive, and the price-mechanics claims are retired here.

Two corollaries follow that the thesis should have drawn earlier:

- **A burn is a buyback.** Economically, retiring supply with revenue is a pro-rata return of value to an existing claim. Buybacks change per-share value; they do not confer moneyness. No quantity of buyback has ever made an equity a store of value, and no burn schedule will do it either.
- **Collateral is a floor, not a premium.** A required lockup denominated in the asset supports a valuation from below. It is a balance-sheet fact, and Section 11.10.9 bounds how large a fact it can be.

The Test That This Reasoning Does Not Prove Too Much

If durable fee extraction from a commoditized flow made an asset money, the world would be full of monetary assets. CME clears standardized futures; DTCC settles equities; Visa takes a spread on card turnover; MSCI licenses indices that everyone benchmarks to; pipelines and port authorities meter physical throughput. All of them sit on flows more commoditized than canonical MatMul, all of them extract fees durably, several are structurally harder to bypass than any protocol will be, and *not one of their equities is money*. Visa’s take rate does not make Visa shares a store of value; it makes them a good business.

The comparison is the discipline. If the argument for the asset’s monetary premium is “recurring fees on indispensable throughput,” the argument is an argument for a high-quality utility equity, and it should be labelled as one. Something else has to be doing the monetary work.

11.10.8 Where a Monetary Premium Would Actually Come From

The constructive half. If fees, burns, and collateral yield a well-valued cash-flow claim rather than money, the monetary claim has to be relocated rather than abandoned — and the place it relocates to is one the thesis has already spent six Parts describing without noticing that it was

the monetary argument.

Δ is not a constant; it is a function of regime pressure. Write $\Delta(R)$, with R the regime-pressure index of Stage A (Section 11.14.1). Section 11.10.5 showed why Δ is thin in benign states: a solvent, suable indemnitor substitutes for a proof. Invert that sentence and it names exactly when Δ is large. The substitute fails when the counterparty:

- **cannot be sued** — cross-jurisdictional, pseudonymous, or sanctioned, so no forum will hear the claim;
- **will not be solvent when it matters** — the failure is correlated, and the indemnity is worth least in exactly the state that triggers it;
- **is the adversary** — a state, a platform, or a regulator, where the entity you would sue is the entity imposing the loss;
- **inflicts a non-compensable loss** — censorship, seizure, deplatforming, or disclosure, where damages do not restore the position because the position was never about money.

That list is the threat model of Chapter 5, restated as a demand curve. The thesis is therefore internally consistent — but the consequence is unflattering to fee accrual as a monetary argument and places the monetary burden on links 8 and 9:

Δ is small in benign states and large in the tail. The asset's usefulness is state-contingent, and its distribution is skewed toward exactly the states everyone else's arrangements stop working in.

State-contingency does two different things, and only one of them is monetary. The tempting move here is to say that an asset whose usefulness rises in states where the marginal utility of wealth is high must be worth more than the discounted expectation of its cash flows, because those cash flows arrive weighted by a high state price. That move fails the test this section's own demolition just set. Section 11.10.7 defined a monetary premium as value *in excess of* discounted cash flows, and a discounted-cash-flow valuation carried out with the correct stochastic discount factor already weights every cash flow by its state price. State-price weighting is what a risk-adjusted discount rate *summarizes* — and a single discount rate summarizes it correctly only when payoffs are not heavily state-contingent, which ours are. A payoff this state-contingent needs the stochastic discount factor itself, not the single scalar that compresses it. An argument that stops there has relocated the monetary claim into the denominator of a DCF, which is not outside the DCF at all. The argument therefore has to be split into two mechanisms, one conceded and one kept.

Mechanism one: a negative beta on the fee stream. Real, valuable, and not monetary. Because $\Delta(R)$ rises when the substitutes fail, the fee stream is countercyclical with respect to regime pressure: the protocol can charge most for its differential properties in the states where

wealth is scarcest. A rational investor pays more per unit of *expected* cash flow for that stream than for a procyclical one, and should. But this is exactly a discount-rate effect. It lowers the cost of capital applied to the same fee stream, it is fully representable inside a discounted-cash-flow valuation, and it therefore produces no moneyiness. A countercyclical utility is a better utility, not a monetary asset. Conceded, and set aside: nothing in the monetary claim below rests on it.

Mechanism two: a holder-side service flow. This is where the moneyiness is. The bearer of the asset obtains something that is not a cash flow at all, and so has no cash flows to discount: the ability to transact, prove, hold, and exit at the moment when the substitutes for proofs, courts, and custodians have stopped working. That service accrues to whoever holds the unit, by virtue of holding it, and not by virtue of any distribution from the protocol. Nothing is paid out, so there is no stream to place in a numerator and no discount rate — correct, risk-adjusted, or otherwise — that reaches it. This is the same category of thing as the value of holding a bearer instrument outside the banking system, and it is why an asset with no cash flows whatsoever can carry a price.

The mechanism is standard, and the analogy is to the mechanism only. That a non-cash-flow service from *holding* a thing can carry a price no cash-flow model reproduces is the storage-theory account of **convenience yield** in consumption commodities (Kaldor, 1939; Working, 1949): the holder of physical inventory obtains an option against stockout that the holder of a claim on future delivery does not. The analogy drawn here is to that mechanism — a service flow accruing to the bearer rather than to a claimant — and not to the setting. A triad asset is not a storage commodity, it faces no stockout risk in the storage-theory sense, and no magnitude in this thesis derives from the commodity literature.

The distinction the mechanism turns on is easiest to see with a fire extinguisher and a contract for the delivery of one. Both have a price. Both can be bought today, valued, resold, and hedged. Only one of them works during the fire. The premium you pay for owning the extinguisher rather than holding a delivery claim on it — the amount by which the physical commands more than the discounted contract — is exactly the quantity no discounted-contract valuation reproduces, and exactly the quantity this thesis is pointing at when it says the holder-side flow has no cash flows to discount. In the moment of use, the contract and the extinguisher are not substitutes at any price, and every argument in this section about regime states is an argument about when the market is on fire.

Gold, which this thesis keeps comparing itself to, is a separate case and should not be filed under storage theory: gold's convenience yield in the storage-theory sense is near zero, which is why its lease rates are negligible and its forward curve is essentially full cost of carry. What is true of gold is the part that matters here. It has no cash flows, negative carry, and industrial

demand far too small to explain its price; every constructible DCF for gold returns a number well below the market. On the standard account that residual is monetary and portfolio demand — what holders pay for a bearer holding that is useful when other arrangements are not — which is the holder-side category described above, arrived at from a different direction.

Regime-Contingent Convenience Yield

The monetary premium available to a triad asset is the **holder-side service flow** generated by the state-contingency of $\Delta(R)$ — the bearer's ability to transact, prove, hold, and exit when the substitutes for proofs, courts, and custodians have stopped working — *conditioned on* credible non-discretion in issuance and on bearer holdability. Those two are not added terms but enabling conditions: the flow is only a monetary service while issuance discipline holds and the unit remains bearer-holdable, and it collapses if either fails. The service accrues by virtue of holding rather than as a distribution, so there is no stream to discount. That is what places it outside a discounted-cash-flow valuation rather than merely at a different discount rate inside one.

It is *not* the countercyclicity of the fee stream. That is a negative beta on cash flows: real, valuable, and captured exactly by a DCF using the correct stochastic discount factor, which is why it confers no moneyness.

It is also not generated by burns, by collateral, or by fee share. Those produce a cash-flow claim and a floor (Section 11.10.7).

Which conditions and which links this actually implicates. The consequence of the split is immediate. If the monetary claim rests on the holder-side flow, then it rests on the properties that determine whether a holder can actually use the thing under pressure — and none of those are fee-routing properties. Of the five lemma conditions, only **issuance discipline** touches the channel at all, and it does so because credible non-discretion is part of the payoff rather than part of the accrual. The other four are conditions for value *accrual*, and accrual is a DCF quantity. Link 8 then requires a persistent self-custodied constituency capable of bearing loss, and the monetary premium hangs off link 9 — liquidity, neutrality, verifiability, legal holdability, and agency preservation (Item 9) — which are exactly the conditions under which a bearer can transact, prove, hold, and exit when it counts. No burn schedule reaches the premium: a burn operates on the claim, and the premium is not on the claim.

Now discount it honestly. The hedge has a defect, and it is the same defect the thesis has already documented without drawing the monetary conclusion. The payoff is negatively correlated with the states you fear — that is the point — but the *delivery* is positively correlated with them. The states in which $\Delta(R)$ spikes are energy interdiction, hardware denial, network filtering, and jurisdictional pressure (Section 5.5, Section 15.10.1), and those are precisely the

states in which the protocol's ability to supply privacy, proofs, and compute is impaired. [Section 31.1.6](#) half-concedes this and Red Line 13 makes the mechanism falsifiable, but both frame it as a threat to the hinge. It is also a *discount on the hedging premium*, which is the monetary consequence, and the honest statement of it is: you are being asked to pay a premium for insurance whose underwriter is exposed to the insured event.

Two further deductions from the same discount:

- **The slots are occupied.** Gold already holds “works when the network does not.” Bitcoin already holds “works when the state does not.” Neither requires a grid at the moment of use — gold requires none at all, and a Bitcoin key survives an outage even where settlement waits. A triad asset needs power, reachable networks, and obtainable reference hardware in order to deliver anything at all, which makes it a hedge against a narrower and more specific band of failure: adversarial conditions severe enough to need proofs and privacy, but not severe enough to take the grid.
- **The band may still be the important one.** That band — administrative repression, surveillance, synthetic media, platform and compute enclosure, financial repression without collapse — is the one Part I argues is most likely. Sovereign optionality ([Section 15.8](#)) exists to widen it, and disruption-adjusted VERIFYPRICE ([Section 15.10.2](#)) exists to price whether the widening is real.

And do not overclaim it. The state-contingency argument establishes a *mechanism*. Non-cash-flow service flows accruing to a bearer are ordinary in asset pricing, and there is nothing exotic in claiming one here. The **holder-side flow is nonetheless unsized**: nothing in this document estimates how large a regime-contingent convenience yield a triad asset could command, and no honest estimate is available before there is a market in which to observe one. This is an empirical question the thesis has not answered. What it can do is refuse to smuggle the answer in — which is why the chain-strength rating for the store-of-value premium is Medium (conditional), why link 8 now asks who bears the asset's loss, and why the relocated claim is attributed to link 9 rather than to fee accrual. The measurement instrument exists so that the question can be settled rather than asserted: [Section 24.8.6](#) publishes the lending-rate, basis, and wrapper-basis series against the regime-pressure index, and a series flat across regime states is the null result, reported as such.

11.10.9 Collateral: Unit-Elasticity, Floor, and Wrong-Way Risk

Condition 3 requires operators to post the native asset as collateral, and the reference designs specify the requirement as a share of *capacity value*: “10–20% of capacity value” in [Section 7.4](#), a “10–15% collateral requirement” in Design A, and “proportional to their capacity” in [Section 11.9](#). That specification has a consequence the document has never drawn.

Collateral Unit-Elasticity

Let K be the USD value of deployed capacity, χ the required collateral ratio, and P the asset price. The number of units that must be locked is

$$N = \frac{\chi K}{P}.$$

The requirement fixes a *value*, not a quantity. Collateral demand is therefore unit-elastic in price: a doubling of P halves the units required to satisfy the same obligation.

(χ is used here for the collateral ratio because κ already denotes net mechanical gain in [Appendix H](#) and facility capacity share in [Section 15.8](#).)

It is a level effect on free float, not a slope on demand. The mechanism immobilizes a fixed *dollar* quantity of float — a level effect on the free float, not a buy order of any kind, since a lockup places no bid — and the valuation it supports grows only as deployed capacity K grows, not as price grows. Describing it as “structural demand that scales with the network” is true only in the sense that it scales with physical build-out, and physical build-out is slow, capital-rationed, and observable.

And K deflates too. Capacity value is the replacement cost of deployed infrastructure. It therefore inherits the same deflation exposure as the fee base in [Section 11.10.2](#): the same generation of throughput is worth fewer dollars of replacement capital each year. A collateral requirement pinned to K shrinks in dollar terms unless deployed physical capacity grows faster than its unit cost falls. This is the same empirical bet, entering through a second door.

Reflexivity runs in both directions, and the second one is not in the document. For small moves the mechanism is genuine negative feedback, and it is a real design achievement: price rises, fewer units are needed to cover the same χK , surplus collateral releases into float and meets the buying; price falls, operators must post more units, and they buy. An automatic stabilizer.

For large moves it inverts. Collateral denominated in the asset it insures loses value exactly when slashing risk and operator distress rise — textbook wrong-way risk, and in its dynamic form the familiar margin-spiral or leverage-cycle mechanism. The operator’s instruction in a drawdown is “post more units, at a moment when your existing holdings are worth less and your operating business is stressed.” A common answer is not to post; it is to *exit capacity*. Exiting reduces K , which reduces required collateral, which releases locked units into a falling market. The stabilizer becomes an accelerant, and it does so through the same equation that

made it a stabilizer:

$$P \downarrow \Rightarrow N \uparrow \Rightarrow \text{operator exit} \Rightarrow K \downarrow \Rightarrow N \downarrow \Rightarrow \text{float} \uparrow \Rightarrow P \downarrow.$$

Nothing in the value-leakage mitigations of [Section 11.9](#) contemplates this loop, and no red line covers it. Naming it is the minimum; a metric for it — collateral coverage against realized slashing exposure, and capacity withdrawal rates conditioned on drawdown — belongs on the Value Capture Board ([Section 24.9](#)) in a later pass, and this document should not pretend it is there yet.

The design tension, stated plainly. Collateral must be the native asset in order to create native demand, which is the whole of Condition 3’s monetary contribution. Collateral should be uncorrelated with the risk it insures in order to be credible, which is the whole of prudential collateral design. **You cannot have both.** Any move toward external collateral — stablecoins, bonded fiat, over-collateralization in an uncorrelated asset — strengthens safety and weakens Condition 3’s monetary contribution by exactly as much. This is not a problem to be solved; it is a trade-off to be chosen deliberately and disclosed.

The magnitude bound. χK is bounded by an accounting identity, and the arithmetic is worth doing even roughly. The following is an order-of-magnitude illustration, not an estimate: the inputs are chosen to be generous rather than defensible, and no claim rests on their precision.

Deployed capacity value K	Locked at $\chi = 0.15$	Comparable
\$1T (deliberately generous)	\$150B	A single large-cap equity
\$100B (more defensible)	\$15B	A mid-cap equity

Table 11.2: Order-of-magnitude ceiling on collateral lockup. Illustrative arithmetic on assumed capacity values, not a forecast; the point is the exponent, not the digits.

Those are large-cap-equity numbers. They are not monetary-aggregate numbers, and no choice of χ inside a plausible range changes the exponent. Collateral therefore cannot be the source of moneyiness. At best it is a floor.

The one real contribution, stated in the document’s own notation. Collateral does do something, and it is worth being precise about what. It creates a holding constituency whose supply is highly price-inelastic — operators cannot sell locked units at any price without exiting the business — which shrinks effective free float and steepens the price response to a given flow. In the Stage C decomposition ([Section 11.14.1](#)), that is an operation on the liquidity and market-impact state Λ_t , not on the native monetary impulse θ_t .

Collateral is a slope effect on price impact, not a level effect on price. By immobilizing float (a level effect on *free float*), it changes how much price moves per unit of flow. It does not establish where price belongs, because it places no bid.

And a slope effect cuts both ways: a steeper response to buying is a steeper response to selling. Calling this “structural demand” overstates it in one direction while concealing the symmetry in the other.

11.11 From Utility Demand to Monetary Premium

Many indispensable services—electricity, bandwidth, compute instances, legal services, cloud storage—have recurring demand without being stores of value. The monetary argument requires explaining why *this* utility becomes **monetized savings demand**, not just operating expense.

The bridge has eight conditions. The middle five are the [Value Capture Lemma](#) arriving early and informally; they are named here in one line each and stated formally at [Lemma 11.1](#), not argued twice:

1. **Utility demand is not enough.** Recurring demand for a service does not automatically create a scarce, holdable asset. Bandwidth is indispensable, but claims on bandwidth do not become money.
2. **The asset must be required for fees or settlement** — required fee medium, Condition 1.
3. **Fees must produce burns, retirement, or staking yield** — supply reduction, Condition 2.
4. **Operators must post the asset as collateral** — Condition 3, whose worth [Section 11.10.9](#) bounds as a floor and a slope effect, not a source of premium.
5. **Issuance must be capped or capacity-constrained** — issuance discipline, Condition 4.
6. **Users must not be able to bypass the asset at equal service quality** — non-bypassability, Condition 5.
7. **Telemetry must prove the loop is working.** Fee coverage, burn rates, collateral lockups, and workload demand must be publicly verifiable—not asserted.
8. **If any value-capture condition fails, the asset does not even capture its own service value.** It becomes infrastructure exposure whose economics accrue to operators or customers rather than to holders. Meeting those five conditions makes service value accrue to the asset; it does not make the asset money, establish stress-deliverable service, or create a loss-bearing holder constituency.

This section is the “utility-token trap” defense. The thesis does not claim that useful things automatically become money; it claims that *under specific, testable value-capture conditions*, a useful asset accrues the value of its usefulness to its holders — a cash-flow claim and a collateral floor. Whether anything further, a monetary premium in excess of those, is available is argued

separately and on different grounds in Section 11.10.8. The Value Capture Lemma formalizes the accrual conditions; the telemetry regime (Chapter 23) makes them falsifiable.

11.12 Why Users Cannot Simply Bypass the Asset

The strongest economic objection to the thesis is simple: *“Why can’t AWS, a ZK prover marketplace, or a privacy wallet sell the same service for fiat or stablecoins and bypass the base asset entirely?”*

The monetary object at stake in this objection is the **base asset** — the fee, settlement, and collateral unit — not Work Credits. Work Credits are bypassable service claims by design, and their bypass threatens only service-claim economics, not the monetary thesis. What the thesis needs is that the *base asset* cannot be bypassed, and the objection is valid unless the protocol enforces native value capture. The required mitigations:

1. **Core fees are denominated in the base token.** Provers, routers, and settlement corridors accept only the native asset for protocol-level fees. Fiat or stablecoin payment requires acquiring the asset first.
2. **Collateral must be posted in the base token.** Operators cannot participate without holding significant quantities of the asset. Calling this “structural demand” overstates it in one direction while concealing the symmetry in the other (Section 11.10.9): it immobilizes a fixed *dollar* quantity of float, which steepens the price response to flow in both directions rather than establishing where price belongs.
3. **Burns create scarcity tied to usage.** A meaningful share of fees is permanently destroyed. Higher triad usage means lower circulating supply.
4. **SLA priority requires the asset.** Gold-tier SLAs, governance participation, and priority access during congestion require holding or staking the asset.
5. **Settlement paths are natively denominated.** Privacy corridors and settlement rails denominate in the base token; off-ramps exist but are not the default.

Honest admission: If users can in practice bypass the asset—for example, if most operators accept stablecoins and immediately off-ramp, or if hyperscaler-hosted provers dominate the market and set prices in fiat—then the SoV thesis fails. The system may remain useful infrastructure, but the monetary claim collapses. The telemetry must track this: native-asset fee share, bypass channel volume, and operator off-ramp rates should be visible on the Economic Coverage Board (Chapter 24).

11.12.1 The Incumbent Bypass Channel Is Not Hypothetical

The five mitigations above answer a proposal. They do not answer a precedent, and a precedent exists: **stablecoins already settle a large fraction of the value the thesis is targeting**, at scale, at near-zero cost, with no native monetary asset anywhere in the loop. Tron and

Ethereum carry stablecoin transfer volumes that exceed most national payment systems; in several emerging-market currency crises, dollar-stablecoin wallets — not Bitcoin, not Zcash — were the observed instrument of capital flight and censorship-adjacent settlement. That is a decade of revealed preference in exactly the states of the world this thesis calls State 3 (Section 27.2), and it is not engaged anywhere in the design.

The stablecoin channel is a live test of Condition 5 that has already been run, and the honest reading is unfavorable in three of four respects:

- **Where stablecoins beat the triad: benign and mid-band repression.** A stablecoin settles in seconds, costs cents, has no learning curve, and — the uncomfortable part — offers *practical* censorship resistance to the median user, because freezing a Tron address is harder than freezing a bank account even though it is easier than censoring a shielded pool. For payments-motivated flight, the bearer-asset properties this thesis prices (privacy by default, portable proofs, verified compute access) are mostly not on the buyer's list.
- **Where the triad can still differ: the upper band.** The stablecoin's weakness is structural and appears exactly where the thesis's services begin: it is a claim on a custodian's reserve and an issuer's goodwill, it freezes under sanctions enforcement, it carries no proof or compute service, and its censorship resistance is borrowed from its host chain rather than owned. As repression crosses from financial into administrative and epistemic forms — compelled disclosure, platform deplatforming, attestation demands — the stablecoin inherits the substitutability of its issuer, not of cryptography. That is the band in which a triad asset's Δ over a stablecoin is nonzero, and Section 11.10.8 already prices the premium as regime-contingent.
- **The uncomfortable arithmetic.** None of that changes the base fact: for the median settlement need in the states most likely to occur, the incumbent bypass channel is cheaper, faster, and already deployed. The thesis's five mitigations make bypass harder *on this protocol*; they do nothing about bypass *around* it. Red Line 6 is therefore best read as already having an empirical prologue: every stablecoin-settled triad-adjacent payment is a data point that the capturable wedge must be measured against, not assumed above.

The design implication is not despair but precision: the asset's fee-bearing services must be ones a stablecoin cannot express — private settlement where the stablecoin is a custodial IOU, proofs where the stablecoin is silent, verified compute where the stablecoin has nothing to verify. Where the triad service is interchangeable with a stablecoin payment, Condition 5 should be presumed failed, and the Value Capture Board's stablecoin-denominated fee share series (Section 24.9) is the instrument that watches it fail or hold.

Co-option as Bypass

The bypass threat is not always a competing product. It can also be institutional co-option: ETFs, treasury companies, margin loans, and regulated custody that deliver ex-

posure to the asset's price without requiring users to interact with the protocol's privacy, proof, or settlement rails. If the majority of demand is satisfied by custodial wrappers that bypass the fee-burn-collateral loop, the asset's *price* may rise while its *monetary thesis* weakens. The telemetry must distinguish between custodial exposure (which does not exercise value capture) and protocol-native usage (which does).

11.13 Wrapper Dominance Risk

Co-option as bypass is dangerous enough to warrant a named, top-tier risk rather than a sidebar. A protocol asset can succeed as a financial product while failing as money. If ETFs, custodians, treasury companies, broker-dealers, lending desks, and stablecoin wrappers satisfy most demand for exposure, the asset's price may rise even as native settlement, privacy usage, fee burns, collateral lockups, and non-custodial flows stagnate. This is not a contradiction. It is the signature of **wrapper dominance**.

Wrapper dominance is dangerous because it turns a monetary network into a reference price. It creates holders without users, exposure without settlement, and liquidity without sovereignty. The asset can moon while the monetary thesis dies.

Wrapper Dominance Ratio (WDR)

Wrapper dominance has a stock component and a flow component, and they must not be divided into one another: exposure is measured *at* an instant, usage *over* an interval, so a single ratio would scale with the arbitrary length of the measurement window. WDR is therefore reported as a pair.

Stock form — what share of economic exposure sits in custodial or synthetic form, where E_t^{wrapper} and E_t^{native} are exposures at time t :

$$\text{CustodialExposureShare}_t = \frac{E_t^{\text{wrapper}}}{E_t^{\text{wrapper}} + E_t^{\text{native}}}$$

Flow form — how activity divides over a stated horizon h , where $V_{t-h:t}$ is volume across that window:

$$\text{WrapperActivityRatio}_{t,h} = \frac{V_{t-h:t}^{\text{wrapper}}}{V_{t-h:t}^{\text{native}}}$$

The stock form is bounded in $[0, 1]$. The flow form is not: it divides one volume by another, so it is bounded below by zero and unbounded above, and it is undefined when native volume is zero. Both are window-explicit. A rising stock share with a rising activity ratio indicates the asset may be financializing faster than it is becoming money.

The horizon h must always be published alongside the flow form.

Inputs to track: ETF/wrapper AUM; custodied balances; exchange balances; wrapped token supply; native fee share; non-custodial settlement volume; shielded/private settlement volume; staking/collateral lockups; fee-burn coverage.

Why wrappers come first, and why that is not reassuring. The observation that first-generation use of a new medium wears the old medium's form (McLuhan 1964; Sources) — early print looked like manuscripts, early television like radio with pictures — supplies the intellectual lineage of this ratio, and it changes what a rising WDR should be read as. It is not, on its own, evidence of pathology: during a medium transition, consumption of the new asset through the old form is the expected initial condition, not an anomaly. What it means is that the WDR is measuring a *phase*, and that the thesis's claim is specifically about whether the transition completes. Wrapper dominance in year three is the manuscript stage of the bearer asset. Wrapper dominance in year twenty is a medium that never finished forming. The failure gate above exists precisely because the frame does not promise completion: first-generation use is always the old content, and the red lines, not the analogy, decide when waiting has become failing.

Gold ran this experiment already. The financial lineage is a century old, and it is worth naming because it removes any temptation to treat wrapper dominance as a novel crypto pathology. Paper claims on gold grew until they dwarfed the metal: unallocated bullion accounts, gold certificates, and eventually exchange-traded products came to represent exposure-to-gold on a scale that holding-of-gold never matched — the vault holds a fraction of what the paper promises, and the fraction is the finding. A shareholder in a gold ETF owns a claim on a trust that owns gold; a holder of unallocated bullion owns, in the stress-relevant sense, a deposit at whoever holds it. Both are exposure. Neither is a coin in a vault, and on the day it matters — the redemption run, the allocation freeze, the jurisdictional block — the difference between allocated and unallocated is the entire difference. Paper gold did not destroy gold. It demonstrated the exact pattern this ratio watches for: the financial claim outgrowing the bearer object it references, prospering for decades, and leaving the underlying's monetary character unchanged while the exposure stack did the volume. That is the precedent the WDR tracks, and the reason it is a stock *and* a flow rather than a single number.

Failure gate: If WDR rises for multiple quarters while native fee share, private settlement, and collateral lockups stagnate, the asset may remain investable but its **SoV-as-protocol** thesis is weakening. This gate is added to the Red Lines in [Section 28.6](#) and tracked on a dedicated Wrapper Dominance Board ([Section 24.9](#)).

11.14 The Market Realization Plane

The Wrapper Dominance Ratio is a **pair** — a stock share and a flow ratio — reported side by side and never divided into one another. The stock component measures how much ownership has migrated into custodial and synthetic form; it does not explain the *flow* mechanics by which those wrappers set prices. That requires one more construct.

Market realization is deliberately **not** an eighth layer of the stack. Layers 0–6 produce privacy, proofs, compute, settlement, and governance. Market realization produces none of those; it is the conventional financial machinery that represents claims on the resulting monetary object. It is orthogonal to the stack, so we model it as a *plane* around the stack rather than a layer within it.

Market Realization Plane

The set of external institutions and instruments through which claims on the native monetary object are represented, financed, allocated, and priced: exchanges and custody, spot ETFs, corporate treasury vehicles, index products, options and futures, leveraged and inverse ETPs, dealer swaps, prime-broker financing, passive mandates, systematic trend strategies, and rules-based or agentic treasury systems.

The result is two nested loops, and the thesis needs both to be instrumented.

The inner protocol loop determines whether the stack works:

Create/Compute → Prove → Settle → Verify

The outer market-realization loop determines how the stack is financially represented:

Narrative → Wrapper → Allocate/Lever → Dealer Hedge → Price → Narrative

The inner loop determines native monetary function. The outer loop determines market-price realization. Either loop can strengthen while the other weakens, and the outer loop can run for years on narrative alone.

11.14.1 A Three-Stage Model

The combined thesis can now be stated as three linked but separable processes.

Stage A — regime pressure creates structural need. Let R_t index financial repression, surveillance, synthetic media, and AI concentration. Compute demand is deliberately *not* a component of R_t : it is a demand driver in its own right, entering alongside regime pressure rather than as part of the pressure index, so that $D_t^{\text{triad}} = g(R_t)$ is not partially defined in terms

of its own output. Regime pressure generates demand for triad capacity:

$$D_t^{\text{triad}} = g(R_t)$$

This is the causal layer developed in Part I.

Stage B — protocol design converts need into native value. Let $\phi_t \in [0, 1]$ measure the effectiveness of the monetary design: required fees, burn share, collateral, issuance discipline, and non-bypassability. Then native captured demand is

$$D_t^{\text{native}} = \phi_t D_t^{\text{triad}}$$

As $\phi_t \rightarrow 0$, the system can remain enormously useful while the monetary object fails. This is *not* the [Value Capture Lemma](#) restated as a coefficient: the lemma is an AND-gate on five conditions, while ϕ_t is an **empirical capture fraction** those conditions constrain. The lemma says partial capture is possible only when the gate is open; ϕ_t measures how much of the triad demand actually flows through once it is.

Stage C — market structure realizes native value as a price path. Observed price depends not only on native demand but on aggregate exposure flows and the liquidity available to absorb them:

$$\Delta \ln P_t = \theta_t + I(Q_t, \Lambda_t) + \eta_t$$

where θ_t is the non-flow fundamental or native monetary drift — defined to *exclude* flow effects, so that all flow impact enters once, through the $I(Q_t, \Lambda_t)$ term, rather than twice — Q_t is aggregate exposure demand (a flow over the interval $[t - 1, t]$), Λ_t is the liquidity and market-impact state, and η_t is residual. Aggregate demand decomposes by source:

$$Q_t = Q_t^{\text{native}} + Q_t^{\text{allocation}} + Q_t^{\text{levered}} + Q_t^{\text{dealer}} + Q_t^{\text{trend}}$$

Only the first term is monetary evidence. The remaining four can dominate the price for extended periods, in either direction. [Appendix H](#) develops the flow mechanics formally — holder flow elasticity, the net mechanical gain coefficient, the recycling boundary, market impact, and volatility drag — and [Section 24.8](#) defines the telemetry that makes the plane observable.

Why a protocol should care about external market structure. A designer might object that this is someone else's problem. It is not. External wrappers affect treasury behavior, collateral demand, governance concentration, liquidity-provider incentives, public understanding, security-budget expectations, regulatory pressure, and the price that the protocol's own participants use when making decisions. The protocol should not manage its price. It must nonetheless instrument the structures through which its monetary claims are represented, for the same reason it publishes VerifyPrice: unmeasured dependencies are where theses go to die

quietly.

11.14.2 The Numeraire Is Not Fixed

Everything above measures the price path in units of account whose own value is assumed constant. That assumption is doing more work than it appears to.

Corollary 11.1 establishes that price is not evidence of monetary adoption. The point here is adjacent and independent: even a price series that *is* informative about the numerator can mislead through the denominator. A rising quotation may report an appreciating asset, a depreciating unit of account, or any combination of the two, and the series alone cannot distinguish them.

Numeraire-Dependence

The property that a measured price path is jointly determined by the asset and by the unit in which it is quoted, such that the same series can support opposite conclusions about real command over resources depending on the denominator chosen.

For this thesis the consequence is a measurement requirement rather than a philosophical observation. The Market Realization Plane exists to prevent financialization from being narrated as adoption. A plane instrumented only in the sovereign unit of account is blind to the case where wrapper-led appreciation and currency depreciation move together—which is precisely the regime [Section 3.1](#) argues is likely.

Requirement. The core market-realization series—the Wrapper Dominance Ratio ([Section 11.13](#)), the Wrapper–Native Growth Gap, and the native fee, burn, collateral, and settlement series—should be published in at least one non-fiat numeraire alongside the fiat series. A reference basket of energy, compute, shelter, and gold is sufficient and has the advantage that each component is a thing the holder might actually need. Where the two numeraires disagree in sign, the disagreement is the finding and should be surfaced rather than reconciled.

A sharper definition of the objective. The thesis has so far spoken of purchasing power, which inherits the numeraire problem it is trying to escape. The more precise statement of what a store of value is for:

Wealth is durable, transferable command over necessary capacity under adversarial conditions.

Each term is load-bearing. *Durable* excludes claims that survive only while an issuer chooses to honor them. *Transferable* excludes capacity that cannot be moved or bequeathed. *Necessary capacity* names the denominator explicitly—energy, shelter, compute, mobility, time—rather

than leaving it implicit in a currency. *Under adversarial conditions* is the requirement of [Chapter 4](#) and the reason [Item 9](#) is a design constraint rather than a preference. This is the definition the rest of the document should be read against.

Worked example: structure does not determine behavior. A useful illustration, because it is easy to get backwards. An asset can be genuinely bearer-oriented, scarce, censorship-resistant, and independently settleable at the protocol layer, and simultaneously trade as leveraged technology beta at the wrapper layer—held through exchange-traded products, pledged as collateral, sized by volatility-targeting mandates, and sold first in a de-grossing episode. Both descriptions are accurate. They describe different planes.

The error in either direction is the same error. Concluding from correlated selloffs that the protocol properties are illusory mistakes the plane for the stack. Concluding from the protocol properties that the price will behave defensively mistakes the stack for the plane. [Appendix H](#) supplies the machinery for keeping them apart, and [Section 24.8](#) the telemetry.

A related precision, since the language is routinely abused: a proof-of-work asset is *energy-linked*, not *energy-backed*. Its issuance and security consume energy; it is not redeemable for a fixed quantity of electricity, and no holder has a claim on a joule. The distinction matters because a genuine claim on capacity—a generation asset, a storage system, a Facility Capacity Receipt—has a payoff structure that an energy-linked bearer asset does not, and conflating them overstates the latter’s hedging properties.

Technology success is not incumbent success. One further reading of the same distinction, which the thesis has stated elsewhere in monetary terms and which generalizes. A technology can transform an economy while destroying the equity, debt, and tokens of everyone who financed its first build-out; railways and long-haul fiber are the standard cases. This is [Corollary 11.1](#) pointed at infrastructure rather than at protocols: the social return and the return to the first capital structure are different quantities, and the second is the one an investor actually receives.

Chapter 12

Service-Instrument Design Space

Here is the question this chapter answers: *once work can be verified, what should be built on top of it?* The triad produces three capacities, and Work Credits turn capacity into claims. But a claim is a design decision, not an inevitability, and the decision has consequences that outlive whoever made it. Pick the wrong instrument for a real buyer, and the buyer quietly routes around you: a treasury that needs a two-year hedge will not buy a perpetual index, and an AI lab hedging inference cost will not buy privacy credits it cannot redeem. The failure mode is not loud. Nobody announces they are leaving; the service just never gets the volume, and the instruments sit on the shelf with impeccable specifications and no counterparty.

That is why this chapter walks the design space rather than picking from it. Monetary candidacy remains confined to the base asset and subject to the Triad Coherence Test ([Section 7.3](#)); what follows is the inventory of everything else you can build, organized by which triad leg the claim draws on, and by who — concretely — would ever hold it.

The legacy labels are kept because they name real demand clusters, but read them as *lenses*, not products. “ZK Money” is not a coin; it is the observation that privacy and proofs are usually bought together. “AI Money” is not a token; it is the observation that compute is bought forward, in size, by parties who fear shortage. Each lens describes a buyer, and each buyer implies a different instrument.

12.1 The “ZK Money” Analytical Lens

The legacy “ZK Money” label groups service instruments that primarily reference **Privacy + Proofs**. In practice this means claims on shielded settlement capacity — corridor LP positions, privacy-rail credits — and claims on zero-knowledge proof capacity such as SNARK/STARK service tiers.

What these share is a particular pairing: bearer-like, censorship-resistant settlement rights, bundled with attestations of compliant behavior. Neither half is worth much alone. Settlement rights without attestations are an offshore account with better cryptography; attestations without settlement rights are an audit trail nobody can act on. The value is in the bundle, which is why regulated actors — who need privacy from counterparties and provability toward supervisors *simultaneously* — are the natural buyer.

One transaction to fix the idea. A corporate treasurer operating under an emergent capital-control regime needs to move payroll to overseas staff without exposing vendor rates to competitors or the transaction graph to the local banking authority, and must still satisfy the parent company’s auditors at quarter-end. She buys corridor credits that settle shielded, and each settlement emits a PIDL receipt her auditors can verify without a viewing key. That single purchase — privacy toward the state, disclosure toward the auditor, both on purpose — is what “ZK Money” was always pointing at. Treasuries under repression, regulated actors needing simultaneous privacy and compliance, and individuals wanting cash-like digital assets are variations on the same shape.

12.2 The “AI Money” Analytical Lens

The legacy “AI Money” label groups service instruments that primarily reference **Compute + Proofs**: claims on verified FLOPs, futures on model-specific inference capacity with proof guarantees, and equity-like positions in proof factories whose output is standardized, verified compute.

The common thread is that compute is the one triad leg bought *forward*. Nobody pre-pays for privacy next year; plenty of parties pre-pay for inference next year. That asymmetry is the entire commercial logic of the lens: it exists because shortage risk, not usage, is what generates a term structure.

One transaction. An AI lab training a frontier model has a contractual delivery window for a customer deployment and has watched spot GPU prices move 40% in a quarter. It buys futures on verified inference for the specific model hash it will serve, with proofs of correct execution attached, locking both price and correctness for the delivery window. The hedge is only as good as the verification — an unverified forward on compute is just a promise from whoever sold it — which is why the claims are typed to workload and proof tier rather than to “compute” in the abstract. AI labs and application builders hedging compute shortage, and investors who believe AI demand outlives any particular model cycle, are the same buyer at different horizons.

12.3 Hybrid Instruments

Hybrids reference all three legs at once, and they exist because real obligations rarely arrive one leg at a time.

A **triad basket** is an index-like instrument backed by diversified Work Credits across privacy, proof, and compute workloads — for an allocator who wants exposure to the stack’s aggregate throughput without forming a view on which leg tightens first. A **corridor-linked compute claim** couples private settlement rights with AI inference capacity: the holder can

pay for and receive verified inference without revealing who is asking, which is the shape an agent acting for a pseudonymous principal actually needs. An **agent service treasury** takes the same idea and makes the holder the agent — an on-chain system holding combinations of typed proof, settlement, and compute claims to fund its own operation, buying its own compute the way a factory buys its own electricity.

One transaction, and the one worth watching. A software agent under contract to serve a customer's inference workload for a quarter holds its own treasury of corridor-linked compute claims, draws them down as it serves, and settles with its customer over privacy rails — no human in the loop, and every drawdown receipted. If that transaction ever runs at scale, the instruments holding it up are hybrids by construction.

The key design constraint is always the same, and it is the chapter's answer to the question it opened with: **claims must stay tightly coupled to verifiable work and capacity**, not to governance mood or marketing narrative. The moment a claim's value can be moved by a vote or a slogan, it has stopped being a service instrument and started being a governance lottery — and the buyer who needed the hedge is already gone.

Chapter 13

The ZK + AI Service Economy

ZK is the accounting engine of this stack. If privacy is the right to speak softly and compute is the ability to think loudly, zero-knowledge proofs are the receipts that make both negotiable. They turn “trust me” into “verify me” without exposing the underlying state. A world that routes value and decisions through proofs rather than paperwork is, in effect, a **zk-economy**.

A zk-economy is not just “more proofs.” It is day-to-day life running on attestations: every payment, inference, bridge, and audit backed by succinct evidence that any party can check.

13.1 Why Open Hardware Is a Precondition for the ZK-Economy

If the prover can cheat, the proof is theater.

Today’s largest ZK systems run on stacks that are almost entirely closed: proprietary GPUs, opaque microcode, black-box TEEs, firmware that can be updated silently overnight. In that world, a “fast prover” with a hidden trapdoor can mint convincing bogus proofs. A government-mandated “secure enclave” with a secret backdoor can exfiltrate witnesses from supposedly private circuits.

You get math-flavored trust, not actual trust.

A genuine zk-economy therefore begins one layer below the circuits, with **verifiable machines** (Layer 0). At least part of the proving path must be open from the RTL up through the software stack.

We will never get perfect certainty about every chip in circulation, but we can move to a regime where:

“This proof came from this class of machine, built from this design, and it would have been extraordinarily expensive to tamper with it without being caught.”

That is enough to make “trust the hardware” a falsifiable claim instead of a sacrament.

Concretely, a zk-economy that wants to support the service markets described by the “ZK Money” and “AI Money” lenses needs three things from Layer 0:

- **Open proving paths:** For canonical circuits, at least one proving stack must be open from RTL through firmware and prover binaries.

- **Sampled, attestable devices:** Lots can be sampled with verifiable randomness, imaged, and subjected to structured tests.
- **Hardware-level SLOs for proof honesty:** Profiles with weak sampling or poor audit history should be priced differently.

13.2 How This Opens a ZK + AI Economy

The same logic extends to AI. If proofs are the receipts of the digital order, AI is the industrial plant that consumes energy and data and emits capabilities.

The zk + AI economy is built on three interacting strata:

1. **AI blockchains (“model chains”):** Duplex-style and Ambient-style PoUW systems treat model training and inference as work functions. Work Credits minted against these workloads remain typed verified-compute service claims.
2. **ZK blockchains (“proof chains”):** Nockchain-style zk-PoW systems coordinate global prover markets and issue typed claims on future proof capacity.
3. **Privacy rails as connective tissue:** BTC↔XMR/ZEC swaps, shielded pools, and private rollups let capital move without custody.

The daily-life version looks almost boring:

- Your phone proves you paid a toll **without revealing your route**.
- A DAO buys inference on a model chain; a proof chain clears the resulting proofs; payment moves over privacy rails.
- An exchange settles cross-chain obligations with batches of zk-proof-backed settlement receipts.
- An insurer prices climate risk from sensor networks whose readings are baked into proofs from open hardware profiles.

Underneath the surface, three demand curves reinforce one another:

1. **Proof demand grows** as systems move from “trusted unless flagged” to “untrusted unless proven.”
2. **Verified compute demand grows** as AI saturates workflows.
3. **Privacy demand grows** as financial repression and surveillance tighten.

We can view these as three economic flywheels that, if the thesis is right, will spin up over time:

Proof flywheel. More systems demand proofs → prover markets deepen → VerifyPrice falls or stabilizes → more systems can afford to demand proofs.

Compute flywheel. More AI in workflows → more willingness to pay for verified compute → more Work Credits minted against AI workloads → more capital available to fund model chains.

Privacy flywheel. More repression and surveillance → more demand for lawful private rails

→ deeper anonymity sets and corridor liquidity → cheaper and safer private settlement → more users adopt privacy by default.

Open hardware is what keeps these flywheels from collapsing into a handful of “trust me” platforms. zk-PoW is what coordinates proof work into a measurable, meterable commodity. AI-PoUW is what turns model work into an asset class whose revenues are natively tied to **triad usage**. Privacy rails are what keep capital flowing between them without recentering custody.

Put in one line:

Open hardware gives us honest machines. ZK turns that honesty into portable guarantees. AI gives us something valuable to spend that honesty on. zk-PoW and AI-PoUW are the mechanisms that weld the three into an economy rather than a collection of clever demos.

Chapter 14

SoV Evaluation Framework

Later, Part V will present an operator/investor checklist that spans the whole stack. Here we sketch the **monetary lens**: how to evaluate whether the base asset is a credible store-of-value candidate while keeping the service and evidence instruments beneath it analytically separate.

Understand what you are doing when you run it. You are not scoring a product; you are auditing a supply chain that runs in two directions at once — upward from machines into money, and outward from one balance sheet to whoever is standing behind it. Every question below is a way of asking one of two things: *can this thing deliver what it promises under the conditions that make it valuable?* and *who absorbs the loss when it cannot?* The first is the question a buyer asks; the second is the question an underwriter asks. A store of value has to survive both, which is why the checklist reads as though written by two people.

One distinction does more work than any other here, so it belongs at the front rather than buried as a note: **legibility is not resilience**. A facility can be perfectly transparent about being catastrophically fragile — publish every receipt, pass every audit, display its full telemetry, and still fail on the first stressed day. The reverse is rarer but real: a robust operation whose interior is opaque. Every measurement regime, this thesis's included, is biased toward the first thing it can see; dashboards reward the visible. So the checklist deliberately separates the questions that ask *can the substrate be seen* from the ones that ask *can what we see withstand anything*, and treats a high score on the first set as establishing nothing at all about the second. Items 10 and 11 exist as two items for exactly this reason. Passing the legibility questions makes an asset *auditable*; only the resilience questions can make it a store of value.

For a candidate base asset X , ask:

1. What triad services does it route value from?

- Exactly which Privacy, Proof, and Compute flows pay fees, burn supply, or require base-asset collateral?
- Are those workloads canonical and well-specified?
- Are the associated Work Credits reported as typed service claims rather than counted as monetary holdings?

2. How is base-asset issuance constrained?

- Is issuance formulaic and transparent, and distinct from Work Credit issuance against DVC?

- Can governance mint outside of those rules? Under what constraints?
3. **Can anyone verify the service linkage?**
 - Are VerifyPrice and workload telemetry public?
 - Can a third party recompute DVC and sample the proofs, settlements, and Work Credit claims that justify fee and collateral activity?
 4. **How does it behave under repression?**
 - If real yields are -300 bps for 3 years, what happens to demand for X ?
 - If on-/off-ramps are throttled, can X still be acquired, held, and spent?
 5. **Is privacy lawful by design?**
 - Are there clean paths for voluntary disclosure that do not reintroduce custody or KYC chokepoints?
 - Are the receipts (PIDL) expressive enough for auditors?
 6. **What is the duration profile?**
 - Does X promise fixed nominal coupons (bond-like) or variable participation in fee flows (equity/commodity-like)?
 - Can a sovereign push its real return negative by fiat?
 - If X is duration-neutral as money, what *separate* instrument, if any, is supposed to warehouse the construction interval of the physical assets the stack requires? Conflating the two is a completeness failure, not a feature (Section 3.3, Section 31.5.1).
 7. **Who can be locked out?**
 - Is access to X gated by one jurisdiction, cloud, or app store?
 - Do Layer-0/1/2 assumptions create hidden chokepoints?
 8. **Can users bypass the asset?**
 - Can users buy equivalent Privacy, Proofs, or Compute service for fiat, stablecoins, or direct cloud contracts without touching the asset?
 - Are fees actually paid in the native asset, or do operators accept and immediately off-ramp alternatives?
 - Are burns and collateral lockups material, or negligible relative to total supply?
 9. **Can institutions legally and operationally hold it?**
 - Is custody support available from regulated custodians?
 - Is the legal classification clear enough for treasury or fund mandates?
 - Is liquidity deep enough for institutional-scale positions?
 - Does governance redirect value away from holders?
 10. **Is the physical infrastructure behind it legible?**
 - Are Facility Capacity Receipts (Section 15.7) available and independently verifiable for a material share of active capacity?
 - Is Physical VerifyPrice (Section 20.8) within published bounds, or does verifying the facility take months of trust-the-vendor diligence?

- Is supply concentrated in a small number of hyperscalers, closed TEEs, or single-jurisdiction facilities (Homestead Ratio / Enclosure Risk Flag, [Section 5.6](#))?

11. Is that infrastructure resilient, or merely visible?

- What is network sovereign optionality, $\mathcal{O}_s^{\text{net}}$ ([Section 15.8](#)), and is it published with its component vector and weights rather than as a bare scalar?
- What share of VERIFYPRICE-tracked capacity sits in jurisdictions with an active curtailment, rationing, or interconnection-denial regime ([Section 5.5](#))? Is Red Line 13 approaching?
- Is the headline verification cost a benign-state figure, or is disruption-adjusted VerifyPrice published alongside it with stated disruption probabilities ([Section 15.10.2](#))?
- Does the resilience premium demonstrably lower disruption-adjusted cost, or is it capital misallocation with a security narrative attached ([Section 15.10.3](#))?

*This is the pair to the previous question, and an operator will try to collapse them. The prior question asks whether the substrate can be *seen*; this one asks whether what we see can *withstand* anything — the distinction opened at the head of this chapter, and the reason Red Lines 10 and 13 are separate.*

12. Is demand wrapped, or protocol-native?

- What share of economic exposure sits in ETFs, treasury companies, and custodial wrappers versus protocol-native settlement (Wrapper Dominance Ratio, [Section 11.13](#))?
- Is price rising while native fee share, private settlement volume, and collateral lockups stagnate?

13. Does it equip holders, or manage them?

- What share of flows are non-custodial and require no global identity ([Item 9](#))?
- Does disclosure stay selective (predicates) rather than converging into a reconstructable dossier ([Section 25.1.2](#))?

A base asset that scores well on these questions remains a **next-generation SoV candidate**:

- It is backed by verifiable work, not by decree.
- Its value is rooted in capacities that a digital civilization must keep buying: privacy, proofs, and compute.
- It is measurable under stress via telemetry, not defended by rhetoric.

14.0.1 Comparison: Traditional Hard Assets vs. the Base-Asset Candidate

For context, it is useful to compare how gold and silver—the canonical “old-world” stores of value—score against the same framework:

Gold and silver excel on **Lindy** (millennia of history), **simplicity** (no software, no network), and **zero counterparty risk in physical form**. They remain rational holdings during the transition window while the triad stack matures. Healthy VerifyPrice, VerifyReach, VerifySettle, DVC, and native-buyer telemetry can establish distinctive digital capabilities; they do not es-

Criterion	Gold / Silver	Triad Base Asset
What supports it?	Geological scarcity; 5,000-year Lindy effect	Conditional holder-side agency plus non-bypassable value accrual from stress-deliverable triad services
Issuance discipline?	Mining (energy + geology)	Constitutional base-asset schedule; separate from DVC-bounded Work Credit issuance
Anyone can verify?	Requires assay or trusted custodian	Laptop-class verification in seconds
Behavior under repression?	Seizure risk at borders; detectable	Shielded settlement; non-custodial corridors
Privacy by design?	Physical possession is private; custody reintroduces counterparty	Default-encrypted; viewing keys for voluntary disclosure
Duration profile?	Zero yield unless lent (counterparty risk)	Duration-neutral base asset; project credit remains separate
Who can be locked out?	Jurisdiction can ban ownership, seize at borders	Layer-0/1/2 resilience; multi-path reachability
Programmability?	None	SLAs, escrow, conditional logic, composability
Settlement speed?	Days (physical); hours (paper with counterparty)	Minutes (privacy corridors)

Table 14.1: SoV comparison: traditional hard assets vs. the conditional triad base-asset candidate. Work Credits are excluded because they are typed service claims.

establish that the base asset is superior money or promote Work Credits into savings instruments.

The thesis does not claim gold is obsolete. It claims only that a base asset may earn a distinct premium for preserving agency across specific digital failure modes—financial repression under surveillance, synthetic media, and compute enclosure—if all nine links hold. The bridge is still load-bearing; the candidate remains under test.

The rest of the thesis will shift back to the **stack angle** (Layers 0–6) and the **telemetry angle** (VerifyPrice/Reach/Settle, DVC, buyer quality, decentralization, governance). The monetary question stays in the background: every design choice is judged by whether the base asset can preserve agency and satisfy the nine-link chain, while Work Credits stay what [Section 11.8](#) established them to be — service instruments — under every regime.

Part III

Infrastructure Layers 0–3

Part I argued that soft guarantees are breaking under debt, AI, and platform control, and set out a threat model plus a seven-layer cypherpunk stack. Part II treated **Privacy, Proofs, and Compute** as distinct services that may support a conditional **base-asset monetary candidate**; Work Credits remain typed, energy- and capacity-anchored service claims rather than claims on a monetary base.

Part III switches to the **infrastructure angle**. It asks a blunt question:

What has to be physically true, and stay true under repression, for the triad to be real?

The answer starts at **Layer 0 – Verifiable Machines & Energy**, then climbs through **Layer 1 – Reachability**, **Layer 2 – Distribution & Execution**, and **Layer 3 – Identity & Claims**. These layers do not directly “store value,” but they determine whether any higher-layer SoV story survives contact with hardware, networks, and identity regimes.

Chapter 15

Layer 0: Verifiable Machines & Energy

Layer 0 is where cryptography stops being metaphor and touches matter.

PoW had an implicit Layer 0: silicon, power, and warehouses hashing in the dark. The trust assumption was: *“ASICs will do what the SHA-256 spec says.”* In practice that meant: *“we trust the vendor, the fab, the firmware, and the power company, and we hope no one has a better ASIC they haven’t told us about.”*

Those assumptions were tolerable when the artifact was a hash nobody outside the system consumed: a compromised fab could win more blocks than it should, but no third party was relying on the output for anything. They stop being tolerable when the puzzle mints **receipts** that other people act on, and when proofs and verified compute become monetary primitives. Giving the work an external buyer raises the Layer 0 requirement rather than lowering it.

15.1 Why Layer 0 Is a Monetary Question

Gold worked as money because geology is hard to fake at scale. Bitcoin worked because hashing cost was hard to fake at scale. In both cases, the monetary story rested on silent assumptions: rocks behave; fabs behave; physics behaves.

The triad inherits those assumptions and tightens them:

- Privacy is only as real as the devices that hold keys and speak on the wire.
- Proofs are only as real as the machines that generate entropy, execute circuits, and sign receipts.
- Compute is only as real as the GPUs/ASICs that claim to have run workloads.

If those machines are opaque, remotely steerable, or quietly biased, **Work Credits degrade into theater:**

- A compromised RNG can turn “unpredictable leader election” into a slow rug-pull.
- A backdoored prover can leak witnesses or mishandle private inputs, even though it cannot emit invalid proofs accepted by a sound verifier unless the proof system or implementation is broken.
- A mandated TEE can become a kill switch for entire clusters of provers and routers.

All of this math still runs on matter.

Every proof, every encrypted wallet, every verified FLOP ultimately lives on a sliver of

doped silicon that almost nobody is allowed to audit. Today’s “trusted hardware” stack is a daisy chain of NDAs: closed-source EDA tools, proprietary IP blocks, opaque PDKs, black-box fabs, sealed packaging, vendor-run attestation services.

Layer 0 Precision

Layer 0 does not make cryptographic proofs sound; soundness comes from the proof system and verifier. Layer 0 makes claims about the physical world credible: which machine produced a receipt, whether inputs were captured honestly, whether randomness was biased, whether witnesses were protected, and whether energy/capacity claims are real. A sound proof system should not require trusting the prover’s hardware for proof correctness. Layer 0 becomes essential when the claim includes physical capture, machine identity, energy use, witness confidentiality, randomness quality, side-channel resistance, or fair participation in useful-work markets.

From a monetary standpoint, Layer 0 asks:

“Can we treat triad capacity as collateral if we don’t know what the machines are really doing?”

The answer is “no.” Triad services delivered by machines we cannot interrogate cannot support a credible base-asset monetary candidate; they are IOUs on a hardware cartel plus the jurisdictions that regulate it.

So the mandate of Layer 0 is:

Translate “trust the vendor” into “trust these verifiable claims about the machine and its power,” or don’t pretend it’s money.

15.2 Design Goals and Non-Goals

Layer 0 has to be ambitious enough to matter and humble enough not to LARP full supply-chain omniscience.

Goals

1. **Verifiability over purity.** We aim for *checkable claims* about machines, not for metaphysical purity. Open RTL where we can; structured sampling where we cannot.
2. **Common knowledge of security.** Different actors should be able to agree on *facts* about hardware profiles, even if they disagree about policy.
3. **Energy anchoring, not energy worship.** Power use should be measurable enough that “Work Credit per joule” is meaningful.
4. **Degradability under attack.** When assumptions fail, the system should *degrade visibly*:

telemetry spikes, profiles are deprecated, Work Credits tied to broken profiles are risk-flagged and may face prospective ineligibility for new collateral uses.

5. **Composable exports.** Layer 0 should emit artifacts that higher layers can consume mechanically.

Non-goals

1. **Perfect trustlessness.** We will not “solve” global hardware and supply chains.
2. **Single-vendor dependence.** Heterogeneity is a feature, not a bug.
3. **Total hardware transparency on day one.** Political and commercial realities exist.
4. **Magical protection against all side-channels.** We assign **budgets** to attack surfaces.

15.3 Hardware as Base Reality for Work Credits

In Part II, Work Credits were defined as **claims on standardized units of triad work** (privacy settlement, proofs, verified compute) anchored to energy and VerifyPrice.

Layer 0 defines the **hardware profile** that each Work Credit type rests on. For a canonical workload W , a **hardware profile** H might specify:

- Chip family and stepping.
- Microarchitectural features (e.g., presence of certain accelerators or TEEs).
- RNG source and test regimen.
- Power metering and thermal envelope.
- Known limitations (e.g., “avoids TEE X due to backdoor Y; uses open core Z instead”).

When a Work Credit of type (W, T) is minted, the receipt can say:

“This unit of work was performed on hardware profile H under conditions C , with proof P and VerifyPrice statistics V .”

Monetarily, that matters because:

- Profiling makes **hardware risk priced** instead of hidden. Credits from “sketchy profile H' ” can trade at a discount.
- It lets different actors pick their risk tolerance: some will only hold Work Credits linked to fully open cores; others will accept mixed profiles in exchange for lower cost or higher performance.

Layer 0’s job is not to tell everyone what risk to take; it is to make the **risk legible and instrumentable**.

15.4 The Layer 0 Feasibility Ladder

A common objection to Layer 0 is: “Open silicon and sampled supply chains sound like moonshots. What can we actually do *this decade* given real-world fabs, opaque GPU stacks, and geopolitical constraints?”

The answer is a **graded trust ladder**. Layer 0 does not require perfection on day one; it requires **measurable progress and falsifiable claims** at each grade. Higher grades provide stronger guarantees; lower grades are acceptable for less sensitive workloads, with telemetry that detects when you’re relying on weaker grades.

Grade	Name	What It Means	Timeline	Trust Residual
L0-A	Best Available Today	Attestation + multi-party audits + reproducible benchmarking + diversity	Now	Trust vendor + third-party auditors; side-channel budget measured but not minimized
L0-B	Sampled & Bounded	Lot sampling + imaging + side-channel budgets + independent lab inspections	1–3 years	Trust sampling methodology; residual risk is unsampled units
L0-C	Partial Open	Open RTL for critical components; open firmware; proprietary accelerators wrapped with proofs	2–5 years	Trust fab + packaging; open logic is inspectable
L0-D	Fully Open	Fully open designs + open PDKs + multi-fab production	5–10+ years	Trust physics + fab process; no single-vendor chokepoint

Table 15.1: Layer 0 Feasibility Ladder

Key principles:

1. **Grades are explicit.** Every hardware profile is tagged with its Layer 0 grade. Work Credits, VerifyPrice dashboards, and SLA tiers reference these grades so users know what trust assumptions they’re accepting.
2. **Telemetry detects reliance on weaker grades.** If 80% of proving capacity is L0-A (vendor-attested) and only 5% is L0-C (partial open), that concentration is visible in dashboards. Users can decide whether to hold Work Credits tied to such a distribution.
3. **Higher grades earn lower risk premiums.** Markets should price Work Credits from L0-D profiles more favorably than L0-A profiles, creating economic gravity toward openness as it becomes available.
4. **Migration paths are first-class.** When L0-B or L0-C options become viable, there are documented procedures to migrate workloads off L0-A profiles without catastrophic downtime.
5. **Failure modes are bounded.** If a specific grade is compromised (e.g., a TEE used in L0-A profiles is broken), the impact is contained to that grade. Higher-grade capacity continues to function; affected Work Credits are risk-flagged, receive collateral haircuts, or face

prospective ineligibility.

15.4.1 Policy Hooks: Economic Treatment by Grade

The following table specifies how each L0 grade affects Work Credit eligibility, pricing, and collateral treatment. These are reference parameters; implementations may adjust within bounds.

Grade	WC Eligibility	Issuance Weight	Collateral Haircut	VerifyPrice Tier
L0-A	All workloads; capped at 60% of total issuance	0.9× (below baseline)	0% (if diversified) to 15% (if concentrated)	Reported separately; flagged if >70% of measurements
L0-B	All workloads; no issuance cap	1.0× (baseline)	0%	Standard reporting
L0-C	All workloads; eligible for “open-profile” premium tier	1.1× (bonus for open hardware investment)	0%; may qualify as “pristine collateral”	Premium tier; highlighted in dashboards
L0-D	All workloads; highest trust tier	1.2×	0%; pristine collateral	Gold tier; used as reference baseline

Table 15.2: Economic treatment by L0 grade.

Interpretation:

- **Issuance weight:** Higher grades earn more WC per unit of verified work, incentivizing investment in open hardware.
- **Collateral haircut:** When WC are used as collateral (e.g., for staking, LP positions, or DeFi), lower grades may face haircuts. L0-A collateral is discounted if a single vendor/TEE dominates.
- **VerifyPrice tier:** Dashboards stratify measurements by L0 grade. If most verification runs on L0-A hardware, this is flagged as concentration risk.
- **Issuance cap (L0-A only):** To prevent over-reliance on vendor-attested hardware, L0-A profiles are capped at 60% of new issuance. Excess work at L0-A earns priority for future slots but not extra tokens.

15.4.2 Quantitative Thresholds (Reference Design)

Read the first row plainly, because the number is doing more work than it appears to. Three-fifths of the world’s proving running on vendor-attested chips is the state the cap forbids: every verification in that world is a signature away from being repudiated by whoever holds the signing key, and the stack’s central promise—that anyone can verify without trusting a platform—quietly reduces to trusting three vendors in sequence. The thresholds exist to make that drift expensive before it is complete. Each is a tripwire that fires while the concentration is still a governance problem, rather than after it has become a settlement failure. Breaching

Metric	Threshold	Consequence if Breached
L0-A share of issuance	>60%	New L0-A issuance paused until ratio falls
Single vendor share within L0-A	>40%	Affected profiles moved to “watch list”; 5% collateral haircut
Sampling coverage for L0-B	<80% of production lots sampled per quarter (at tier minimum rates)	Grade downgraded to L0-A until coverage restored
Open-source audit currency (L0-C/D)	>12 months since last audit	Grade downgraded one level

Table 15.3: Quantitative thresholds for L0 grades.

one does not crash anything; it pauses issuance, raises a haircut, downgrades a grade. The damage is contained by making the drift itself the first thing to pay.

These thresholds are governance parameters, adjustable via protocol upgrade with super-majority. Changes are announced 90 days in advance and visible in dashboards.

15.4.3 Pragmatic Starting Point (L0-A)

Today’s stack can achieve L0-A by combining:

- **Multi-vendor diversity:** No single chip family or TEE dominates more than 40% of critical capacity, aligned with the single-vendor share trigger of [Section 15.4](#).
- **Third-party attestation audits:** Independent labs verify that attestation claims match device behavior.
- **Reproducible firmware and benchmarks:** All prover/router firmware is built reproducibly; benchmark results are publicly verifiable.
- **Side-channel measurement:** Known side-channel leakage is measured and published as a “leakage budget” per profile.
- **Cryptographic agility:** Profiles document which primitives can be upgraded via firmware vs. require hardware swap.

This is not trustless. It is *trust-bounded and measured*. The residual trust is explicit: “We trust these vendors, these auditors, and this sampling methodology, and here is the evidence.”

15.4.4 The Path Forward

- **L0-A → L0-B:** Fund independent lot-sampling programs. Publish inspection results. Build statistical models of detection confidence.
- **L0-B → L0-C:** Invest in open RISC-V cores, open RNG designs, open firmware stacks. Wrap proprietary accelerators with proof interfaces so their internal logic doesn’t need to be trusted.
- **L0-C → L0-D:** Support open PDK initiatives. Diversify fab sources across jurisdictions.

Build ecosystem gravity so open designs become economically competitive.

15.4.5 Why This Prevents “Layer 0 Is Impossible, Therefore Thesis Fails”

The thesis does not require L0-D today. It requires:

1. Clear grading of what trust assumptions each profile carries.
2. Telemetry that makes those assumptions visible.
3. Economic and technical paths toward stronger grades over time.
4. Failure modes that degrade visibly, not silently.

If these four conditions hold, Layer 0 becomes a *progress metric* rather than a moonshot prerequisite. The stack can operate today at L0-A while building toward L0-C/D, and users can decide what risk they’re willing to accept at each stage.

15.5 Concrete Components of Layer 0

Layer 0 is not a single device. It is a **bundle of practices and mechanisms** that give higher layers a surface to stand on.

15.5.1 Open Designs Where Possible

The gold standard is **open cores and toolchains**: open RTL or microarchitectures for CPUs/accelerators, open PDKs where geopolitical conditions allow, and reproducible build infrastructure.

Where open options exist, they are **first-class citizens** in hardware profiles. On the margin, this creates **economic gravity**: as open hardware matures, Work Credits tied to open profiles should command a lower risk premium.

15.5.2 Attested Randomness and Entropy

Randomness is the hidden spine of consensus and proofs. A biased RNG can leak signing keys, predict leader election, and make PoUW “fairness” an illusion.

Layer 0 requires:

- Hardware RNG designs that are **documented and testable**.
- **Attested entropy tests**: periodic statistical test suites.
- **Blended randomness**: mixing hardware entropy with VRFs, commit-reveal, and cross-machine aggregation.

15.5.3 Attestation Without Priesthood

Modern hardware ecosystems push TEEs and attestation as the answer to everything. Used naively, they simply move “trust the vendor” into “trust the vendor’s signing key.”

Layer 0’s posture:

- TEEs and hardware attestation are **useful tools**, not **root of trust**.
- Attestations should be **wrapped and sampled**, not taken as gospel.
- Devices emit **local attestations** wrapped in **SNARKs or STARKs** where possible.

15.5.4 Lot Sampling and Destructive Audits

Because we cannot open every chip, Layer 0 leans on **lot sampling**:

- For each hardware profile, a fraction of units are randomly selected for **deep inspection**: decapping, imaging, side-channel probing.
- Results are published as part of the profile’s dossier.
- A protocol-level **Layer 0 Assurance Fund** (2% of fee revenue) finances sampling.

This is how Layer 0 turns “we hope the vendor is honest” into:

“We have inspected a statistically relevant sample of this profile, and here are the findings and residual risks.”

Sampling Economics and Sufficiency

Critics will ask: “How much sampling is enough? Who pays? What confidence do we actually get?” The honest answer begins by conceding the premise: sampling does not produce certainty, and no number in this subsection should be read as though it does. What sampling produces is a *published, priced, third-party* estimate of how much of a fleet is compromised—an instrument with a long commercial pedigree. Lloyd’s Register sold exactly this product for a century: an “A1” grade was not a promise about any individual hull, it was a sampled attestation, published, paid for by the industry that needed it, and trusted by underwriters because the sampling was independent of the shipowners. Marine insurance priced off that grade the way collateral markets here price off L0 grades. The statistics below are the same instrument with the hulls swapped out for silicon.

Sample rate model (reference design):

Statistical sufficiency: For a profile with N deployed units, sampling n units provides confidence that:

$$P(\text{detect if } \geq k\% \text{ are compromised}) \geq 1 - (1 - k/100)^n$$

This bound holds under a **uniform-prevalence assumption**: it models compromise as

Profile Criticality	Definition	Min Sample Rate	Funding Source
Tier 1 (Critical)	Profiles in the top 5 by WC issuance OR above 10% of total capacity; a profile meeting either condition is Tier 1	≥ 50 units/quarter	Protocol assurance budget (2% of fee revenue)
Tier 2 (Standard)	Profiles with 1–10% of capacity	≥ 20 units/quarter	Protocol assurance budget
Tier 3 (Emerging)	New profiles in probation; $< 1\%$ of capacity	≥ 5 units/quarter	Profile sponsor (operator or vendor)

Table 15.4: Sample rate model by profile criticality.

spread evenly across the deployed population, so that a random sample of n units meets compromised units at the background rate. The assumption is where the bound is weakest. Concentrated compromise—an entire production lot backdoored while others are clean, or 20% of lots attestation-only and never destructively inspected—can defeat the bound, because prevalence in the sampled lots is then not prevalence in the fleet. This is precisely why lot selection is randomized and stratified rather than convenience-sampled, and why the attestation-only residual is a published quantity rather than a silent one.

For Tier 1 profiles ($n = 50$, $k = 5\%$), detection probability is $\sim 92\%$. This is not certainty—it is **bounded uncertainty**, documented and priced.

The converse matters just as much and must be stated with equal care. With n samples and zero findings, the 95% upper confidence bound on the prevalence of compromise is approximately $3/n$ (the rule of three). Fifty clean samples bound undetected compromise at $\sim 6\%$ (the $3/n$ rule; the exact bound $1 - 0.05^{1/50}$ is 5.8%), not at zero. Lot sampling supports *graded confidence*, never exoneration, and no claim in this thesis should be read as the latter.

What “Gold tier” requires: A profile achieves Gold tier (= L0-D, matching the grade-to-tier mapping of Section 15.4: Premium = L0-C, Gold = L0-D) when:

1. $\geq 80\%$ of production *lots* are sampled each quarter at or above the tier minimum rate, with the remaining lots covered by supply-chain attestation.
2. No critical findings (backdoors, RNG bias $>$ threshold) in the last 4 quarters.

3. Side-channel leakage remains within published budget under independent testing.
4. At least 2 independent labs have conducted inspections.

Coverage is counted in *lots*, not deployed units. Inspections are destructive and expensive; the lot is the auditable unit a sampling program can afford to buy. “80% of deployed units” would describe a near-total destructive audit of the fleet, which no assurance budget can fund and no fleet can survive; the reference design does not propose it and this document should not be read as though it does.

Gold tier unlocks:

- Full WC eligibility (no issuance cap).
- Zero collateral haircut.
- Eligibility for “pristine collateral” designation in DeFi integrations.

These unlocks are the table’s own consequences, restated. L0-C carries the same 0% haircut and also qualifies as pristine collateral; L0-B earns the 0% haircut but *not* pristine status, and “pristine collateral” as a designation means L0-C or L0-D throughout this document.

Funding mechanism: The protocol allocates 2% of fee revenue to a **Layer 0 Assurance Fund**. This fund:

- Contracts with independent hardware security labs.
- Publishes RFPs for sampling campaigns.
- Maintains a public registry of inspection results.
- Is governed by a multisig of hardware security researchers (not protocol developers).

Transparency: All sampling results are published within 30 days. Raw data (images, test logs) is archived and available for independent verification. If a lab’s findings are disputed, a second lab can be commissioned for arbitration.

15.5.5 How This Opens New Economies

Verifiable machines don’t just make today’s cloud slightly less sketchy. They unlock whole categories of economic arrangement that aren’t viable when hardware is a black box.

Verifiable cloud and compute co-ops. If you can spin up a rack of open-design TEEs or accelerators and have them produce attestations and ZK receipts that anyone can verify, then:

- A small data center in Nairobi or Reykjavík can sell the same class of “trusted inference” or “trusted proving” as a hyperscaler in Virginia.
- A co-op of households can pool “AI appliances” in their basements and earn by running verified work for others.
- Regulators and enterprises can enforce compliance through proofs and telemetry, not vendor logos.

Compute ceases to be a winner-takes-all brand game and becomes a commodity with open admission.

Civic infrastructure that doesn't depend on one vendor's conscience. Voting machines, digital ID kiosks, public-health dashboards: today they are RFPs to a short list of contractors. With open hardware and proof-wrapped attestation, you can build ballot boxes and ID hardware whose entire stack is open to public inspection and require that each device emit public proofs of correct behavior. Democratic legitimacy becomes a property of math and sampling, not of which vendor's logo sits on the plastic.

Tamper-evident telemetry. As the physical world fills with sensors and actuators, the ability to sell tamper-evident telemetry becomes critical. An environmental sensor built on an open profile, with attested firmware and ZK-wrapped readings, can sell CO₂ or temperature data as evidence into climate, insurance, and industrial hedging markets—not just as numbers on a dashboard.

Hardware-native monetary instruments. Once machines themselves are verifiable actors, we can imagine “miner-notes” backed by the future output of a specific open-design proving farm, or municipal bonds whose coupon is denominated in SLA-backed capacity: X proofs per second, Y verified FLOPs, Z private swaps per month, all attested by open hardware meters. These become new kinds of collateral: claims on capacity the world must keep buying.

Open profiles manufactured at multiple fabs in multiple jurisdictions turn chip supply into a multi-polar fabric instead of a single chokepoint. Neutral money needs neutral hardware; verifiable machines produced on a diversified manufacturing base make that a reachable design goal rather than a slogan.

15.6 Verifiable Power: Energy as First-Class Input

If Work Credits are energy-anchored claims on triad work, then **power** is not just an environmental footnote; it is an input to service delivery. It does not back the base asset or make the credit monetary.

Layer 0 treats power in three ways:

1. **Measurement.** Prover farms track power draw, mapping between workloads and power, and local generation vs. grid intake. These feed into **Work Credit metadata** and **VerifyPrice** models.
2. **Resilience.** Micro-grids or backup generation for critical infrastructure; geographic and jurisdictional dispersion of power sources.

3. **Policy hedging.** Transparent power telemetry allows answering “how much of this is actually training/inference/proof that humans pay for?”

For Work Credits, this means credits can carry **energy provenance tags** and markets can price credits differently based on energy profile.

15.7 Facility Capacity Receipts: Beyond Energy

Facility Energy Receipts (FERs, [Appendix A.2](#)) make energy legible, but verified digital infrastructure depends on more than kWh. A facility can have cheap energy and still be fragile if it sits behind a congested interconnect, depends on a water-constrained cooling system, runs in a jurisdiction with curtailment risk, or relies on a single transformer, fuel source, cloud vendor, or hardware profile.

We therefore extend FERs into **Facility Capacity Receipts**. An FCR records the physical capacity envelope behind a unit of verified work: grid node, interconnection class, outage history, backup duration, cooling dependency, water intensity, transformer redundancy, fuel mix, curtailment exposure, jurisdictional power risk, hardware profile diversity, and physical security status.

Facility Capacity Receipt (FCR)

A signed, auditable receipt that extends Facility Energy Receipt data with infrastructure resilience, grid, cooling, jurisdictional, and hardware-diversity metrics.

Table 15.5: Suggested FCR fields

Field	Why it matters
Grid node / balancing authority	Captures congestion and curtailment risk
Interconnection class	Captures firmness of power access
Outage history	Captures reliability
Backup duration	Captures blackout survivability
Fuel mix	Captures geopolitical/commodity exposure
Cooling dependency	Captures water and heat constraints
Transformer redundancy	Captures single-point grid failure
Hardware profile mix	Captures compute centralization
Jurisdictional risk score	Captures seizure, power rationing, sanctions risk
FCR confidence grade	Allows risk-weighted Work Credit issuance

A Work Credit backed by resilient, diversified, verifiable capacity is economically different from one backed by cheap but brittle compute. The market should be able to price that difference. The stack should not hide physical fragility behind cryptographic elegance. If FCR data is unavailable or unverifiable for a facility, credits from that facility should receive issuance

caps or risk haircuts (Section 23.6.4); Section 20.8 formalizes the cost of auditing FCR claims as **Physical VerifyPrice**.

15.8 Sovereign Optionality: Aggregating FCR Into One Exposure

Table 15.5 gives an allocator ten fields per facility. That is the right raw material and the wrong interface. Nobody underwrites a monetary asset by reading ten columns across a thousand facilities, and a field-by-field presentation makes it easy to be locally reassuring while globally fragile: every facility can look acceptable while the network sits inside two balancing authorities and one jurisdiction.

We therefore define a single aggregate exposure over fields the stack *already collects*. This is deliberately not new telemetry. It imposes no additional reporting burden; it changes how existing FCR data is read.

Sovereign Optionality ($\mathcal{O}_s$)

A capacity-weighted composite of substitutability attributes—fuel diversity, firmness, interconnection, jurisdictional dispersion, and hardware mix—computed from Facility Capacity Receipt fields. It *correlates with* the number of independent physical pathways by which verified work can continue under disruption, but does not literally count them; the pathway-count question is answered by Delivered Verified Capacity (Section 15.9).

15.8.1 Adaptation of Scale

The concept is borrowed from state-level energy strategy, where analysts assess how many independent ways a sovereign can obtain electricity, fuel, and industrial feedstock when ordinary market relationships fail. The transposition to a proving network is close but not exact. Earlier versions dropped industrial and logistical throughput on the ground that proving has no conversion analogue. That claim was wrong.

A proving facility has a measurable conversion chain:

energy/fuel → firm electricity → transformer/switchgear → cooling/water → hardware → network → work

It is not manufacturing throughput, but it is conversion throughput. Any edge can become the binding cut, and a weighted resilience score can conceal a zero-flow series bottleneck. Sovereign Optionality therefore remains the substitutability measure; Delivered Verified Capacity below supplies the missing throughput measure.

Table 15.6: Sovereign optionality components and their FCR sources

Term	Component	FCR field(s) used
D	On-site availability	Fuel mix; local generation vs. grid intake (Section 15.6)
F	Conversion flexibility	Fuel mix diversity: Shannon entropy over independent fuel-source shares, normalized by $\log(n_{\max})$ so $F \in [0, 1]$, matching the other components' $[0, 1]$ normalization
I	Reserve depth	Backup duration
G	Grid firmness	Interconnection class; outage history; transformer redundancy; grid node
V	Supplier & technology diversity	Hardware profile mix
C	Coercion exposure (subtracted)	Jurisdictional risk score; curtailment exposure

15.8.2 Mapping to Existing FCR Fields

Every input already exists. Note in particular that the jurisdictional risk score of Table 15.5 is already specified to capture “seizure, power rationing, sanctions risk.” The stack has been collecting the coercion term all along without aggregating or acting on it.

15.8.3 Facility and Network Form

For a facility f , with each component normalised to $[0, 1]$:

$$\mathcal{O}_s(f) = w_D D_f + w_F F_f + w_I I_f + w_G G_f + w_V V_f - w_C C_f$$

The network-level figure is the capacity-weighted mean across facilities, less a concentration penalty. Let κ_f be facility f 's share of verified capacity. H is the capacity-share Herfindahl–Hirschman index computed over the **union of the jurisdiction-partition and the balancing-authority-partition**: two HHI terms, one per partition, *averaged*. That single convention is pinned here; no other aggregation of the two partitions is used anywhere in this document:

$$\mathcal{O}_s^{\text{net}} = \sum_f \kappa_f \mathcal{O}_s(f) - w_H H, \quad H = \frac{1}{2} (H_{\text{jurisdictions}} + H_{\text{balancing authorities}})$$

The penalty term is what prevents the local-reassurance failure. A network of individually excellent facilities that all sit behind the same interconnect or inside the same legal regime has one pathway, not many, and the weighted mean alone would not say so.

One overlap in this form is deliberate and should be visible rather than silent: jurisdictional exposure enters twice, once per-facility through $w_C C_f$ (siting-level risk: seizure, rationing, sanctions) and once network-level through $w_H H$ (concentration of the fleet inside few jurisdic-

tions or balancing authorities). Because the two terms price different failure modes at different scopes, both are retained; the weights w_C and w_H are therefore *published jointly and reviewed annually*, so the overlap is a documented modelling choice an allocator can see and reweight, not a hidden double count.

Relationship to existing metrics. H overlaps deliberately with the dispersion metrics already used for verification monoculture and enclosure risk (Section 28.6, Red Lines 3 and 11). $\mathcal{O}_s^{\text{net}}$ should consume those existing measurements rather than introduce a parallel concentration statistic. Where they disagree, that disagreement is itself a finding.

15.8.4 Governance of the Weights

The weights w_i are judgement, not measurement, and an index whose weights can be quietly retuned is a marketing instrument. Three constraints follow:

- **Published and versioned.** Weights ship with the telemetry spec; changes are proposals with rationale, not silent parameter updates.
- **Independently reproducible.** A third party holding the same FCR data must be able to recompute $\mathcal{O}_s$ and obtain the same number.
- **Subject to capture review.** Weight-setting falls under the same telemetry-capture constraints as the rest of the scoreboard (Section 28.6, Red Line 4). The party being scored must not set the scoring function.

Report $\mathcal{O}_s$ alongside its component vector, never as a bare scalar. A single number is what an allocator wants and also what conceals which substitutability attribute is missing; publishing both is the only honest form.

15.8.5 What $\mathcal{O}_s$ Is For

The index does real work in two places, and is otherwise decoration:

1. **Risk haircuts.** It extends issuance discipline from “we cannot see this facility” to “we can see it clearly and it is fragile” (Section 23.6.4).
2. **Falsification.** It supplies the measurable condition for the energy-sovereignty red line (Section 28.6).

Absent those two hooks, $\mathcal{O}_s$ would be another dashboard number. With them, it is a constraint that can bind against issuance and a threshold that can retire the thesis.

15.9 Delivered Verified Capacity: Surviving Service Flow

Picture a city’s water system. The reservoir level is the number a brochure quotes; what comes out of the tap is what the city actually drinks. Between the two runs a network of mains,

valves, and pumping stations, and any one of them—not the reservoir—can be the thing that decides whether water arrives. DVC is the tap, not the reservoir. Nameplate capacity says how much could flow if nothing failed; Delivered Verified Capacity says how much does flow once a scenario has removed, throttled, or delayed the edges it removes, throttles, or delays. The narrowest main in the network is the minimum cut, and it is the only number that matters while it binds. This is why the measure that follows is a flow computation rather than a capacity sum: the sum flatters, the cut confines. And it is why substitution latency enters the arithmetic at all—a backup reservoir that takes three weeks to route is not redundancy for a fire that started today.

Delivered Verified Capacity (DVC)

For canonical workload W , interval t , and disruption scenario s , Delivered Verified Capacity is the maximum usable service *flow per unit interval* through the complete energy-to-settlement network after scenario-specific edge capacities, substitution latency, and common-cause failures are applied. It is a rate, not a stock.

DVC is defined here but deliberately parameterized by objects defined later: the workload W is drawn from the canonical workload registry of Layer 4, and the sink is usable *settled* service in the sense of Layer 5 settlement. The same W and the same sink are used throughout this Part and the two economic layers; nothing downstream redefines them.

Let $G_{W,t,s} = (V, E)$ be a directed capacity graph whose source nodes are available energy and fuel and whose sink is usable settled service. Each edge e carries surviving capacity $c_e(W, t, s)$ in a common workload-specific service unit. Then

$$\text{DVC}(W, t, s) = \text{MaxFlow}(G_{W,t,s}, c(W, t, s)),$$

and the corresponding minimum cut identifies the bottleneck that actually limits delivery. The unit is not a generic FLOP: it is a settled, independently verifiable unit of the canonical workload at the promised tier. Because DVC is a flow, the quantity comparable to it is issuance over an interval, and the bound below carries the interval explicitly.

Scenario specificity. A benign-state graph and a drought, fuel-shock, transformer-loss, chip-embargo, network-partition, stablecoin-depeg, or coordinated-policy graph are different objects. Scenario s changes edge capacities and may remove nodes or entire dependency classes. Publishing only the expected or capacity-weighted average is insufficient because a series path fails at its narrowest edge.

Substitution latency. Redundancy exists only when an alternative can enter before the service obligation expires. For each edge publish a substitution time $\ell_e(s)$ and the workload's

maximum tolerated interruption L_W . An alternative with $\ell_e(s) > L_W$ does not contribute to DVC in that scenario, even if it eventually restores capacity. Spare hardware in another jurisdiction is not current redundancy if export approval, installation, synchronization, or key rotation takes longer than the SLA permits.

Common-cause dependencies. Nominally separate facilities may share a transformer vendor, firmware signer, cloud identity layer, cable, gas market, cooling basin, stablecoin, legal safe harbor, or market maker. Assign each edge one or more dependency groups g . A scenario can haircut or remove all edges in a group simultaneously. Geographic dispersion without dependency-group dispersion is not independent capacity.

Relationship to Sovereign Optionality. Sovereign Optionality is a capacity-weighted composite of substitutability attributes that *correlates with* independent pathway count but does not count pathways. DVC asks how much usable service reaches the sink after the scenario binds, and is the measure that actually enumerates flow. A network can have high optionality and low DVC because all alternatives are small or slow; it can have high benign DVC and low optionality because one large path supplies nearly everything. Report both, with the active minimum cut and dependency groups, rather than combining them into a reassuring scalar.

Issuance discipline. Work Credits, capacity vouchers, and any other forward service claim must be bounded by a conservative stress-adjusted quantity,

$$Q_{W,t}^{\text{issue}} \leq \min_{s \in \mathcal{S}_{\text{constitutional}}} \{ \text{DVC}(W, t, s) \} \times h_{W,t}, \quad 0 < h_{W,t} < 1,$$

read as: the quantity issuable over an interval t is at most the scenario-minimum *rate* scaled by the interval length, $Q_{W,t}^{\text{issue}}(t) \leq \min_s (\text{DVC}_s \times |t|) \times h_{W,t}$, where $h_{W,t}$ is a published prudential haircut. The interval factor is what makes the rate-versus-stock comparison dimensionally consistent. $\mathcal{S}_{\text{constitutional}}$ is the pre-declared set of scenarios the constitution treats as binding for issuance; the default is all scenarios with non-zero DVC under the reference stress mix. Scenarios carrying zero flow do not silently vanish from the minimum: they are excluded only if the constitution pre-declares them as force-majeure, and that exclusion list is itself a published governance parameter, subject to the same change-disclosure rules as any other. FCR, FER, and PIDL artifacts supply evidence for edge capacities; they are not themselves capacity. Governance may not replace this bound with nameplate power, installed hardware, gross proof throughput, or a benign-state average.

15.10 Resilience Versus Affordability: The Layer 0 Cost Tension

The previous section asks for resilient, dispersed, fuel-diverse, backup-equipped capacity. The thesis elsewhere insists that verification must stay cheap: “anyone can verify” is the hinge, and Section 28.6 treats its failure as fatal.

These two demands pull against each other, and the tension should be stated rather than glossed.

Resilience is insurance. Insurance has a premium. A jurisdictionally dispersed, fuel-diverse, backup-equipped fleet costs more per verified unit than a single facility sitting on cheap interruptible power.

If hardening Layer 0 raises cost per verified unit, then hardening Layer 0 pushes the stack *toward* the very red line that hardening was supposed to protect. A thesis that demands both cheap verification and expensive redundancy without reconciling them has an open flank.

15.10.1 Two Different Exposures

The tension dissolves partially once we separate two things that “energy cost” conflates.

The verification side. Checking a proof happens on reference hardware, at the edge, in small amounts. Its physical exposure is not bulk industrial power. It is:

- **Reference hardware obtainability.** Export controls, sanctions, or platform lockdown can make the reference verifier unavailable, or available only in attested-and-permissioned form.
- **Edge power and connectivity cost** relative to local income, which is what makes verification affordable *in practice* rather than in principle.

The proving side. Generating proofs and running verified compute consumes bulk power. Its exposure is:

- **Energy price and firmness** at facility scale.
- **Competition for the same electrons** from industrial load and state-directed AI build-outs.
- **Curtailement, rationing, and interconnection denial**, which are policy instruments, not market outcomes.

Conflating these produces sloppy claims in both directions. Bulk electricity prices do not directly break “anyone can verify,” because verification is not a bulk-power activity. What breaks the hinge on the verification side is loss of access to unprivileged reference hardware. What bulk energy conditions threaten is the *proving* economy: Work Credit issuance, geographic concentration, and enclosure.

Scope of “Exogenous”

Section 20.7 describes Physical VerifyPrice SLOs as “constitutional” and “exogenous to token price.” That is correct and worth keeping: the SLOs must not be redefined because the market moved.

But exogenous to token price is not exogenous *simpliciter*. Physical VerifyPrice remains endogenous to the physical and jurisdictional conditions under which reference hardware and power are obtainable. Those conditions are partly set by states. The constitutional claim should be read with that scope, not as a claim that the metric floats free of politics.

15.10.2 Disruption-Adjusted VerifyPrice

For the proving side, comparing a resilient fleet against a brittle one at benign-state spot cost is the wrong comparison. It prices the premium and ignores what the premium buys.

What this subsection prices is therefore **fleet resilience of verification**: $\text{VerifyPrice}_s(W)$ below is the cost of verification for workload W in state s , where s runs over the five price states. Two same-named metrics must be kept apart. *Physical VerifyPrice* (Section 20.8) is the Layer 0 audit cost of checking FCR claims—a measurement of how expensive it is to verify the physical collateral story. The *VerifyPrice* SLOs of Layer 4 are the constitutional cost band on verifying a workload’s receipts. They are different metrics that share a surname, and the decision rule at the end of this section reads on the latter.

The object under measurement is a *hardening fleet*: a collection of proving facilities deliberately built redundant, dispersed, and fuel-diverse so that verification survives the states in which it is most needed. Hardening is not decoration. A fleet that has paid for geographic and jurisdictional dispersion is one whose verification still runs when a single region is curtailed, a single transformer vendor is compromised, or a single jurisdiction is excluded—which is precisely when the cost of checking becomes the difference between a claim being auditable and a claim being whatever the surviving operator says it is.

Let s index physical disruption states—normal operation, regional curtailment, extended outage, hardware supply interruption, jurisdictional exclusion—with probabilities π_s and state-conditional cost $\text{VerifyPrice}_s(W)$ for workload W . Then:

$$\text{VerifyPrice}_{\text{adj}}(W) = \sum_s \pi_s \cdot \text{VerifyPrice}_s(W)$$

Two scope notes are recorded with the definition, and kept brief here. First: the states s are the **five price states** enumerated above, not DVC’s capacity scenarios (Section 15.9)—two distinct enumerations serving distinct purposes, and they must not be merged. Second, on

dimension: the sum is **componentwise over the five-field tuple** $(p_{50,t}, p_{95,t}, p_{50,c}, p_{95,c}, \text{fail})$, each field state-averaged separately; when a scalar headline is quoted for $\text{VerifyPrice}_{\text{adj}}$ it is the $p_{95,c}$ field. (Physical-component tuples such as $(t_{50}, t_{95}, e, m, f)$ are views derived from receipts via a per-run cost model, and cost percentiles are percentiles of per-run cost.)

Disruption-Adjusted VerifyPrice

The probability-weighted cost of verification across physical disruption states, rather than the observed cost under benign conditions. A convenience summary; no red line or acceptance rule reads on it.

A resilient fleet has higher $\text{VerifyPrice}_{\text{normal}}$ and lower $\text{VerifyPrice}_{\text{adj}}$ whenever the π -weighted reduction in disruption-state cost outweighs the π -weighted benign-state premium, i.e. whenever $\sum_{s \neq n} \pi_s \Delta_s > \pi_n \Delta_n$, where Δ_s is the cost change in state s and n denotes normal operation. A brittle fleet optimised purely for spot cost reports an attractive headline number and carries the tail.

The π_s are estimates and should be published as such, with their basis, rather than presented as measurements. An operator who can move π_s by assertion can make any fleet look resilient; disruption-state probabilities are therefore governance parameters subject to the telemetry-capture constraints of Section 28.6. Because operators may set self-flattering π , **cross-fleet comparisons use a common reference vector π^{common} published by the observatory**; an operator's self-published π may be shown alongside for context but is never used for cross-fleet ranking.

One further limit belongs beside the definition, because the expectation form invites exactly the misreading it should prevent. $\sum_s \pi_s \text{VerifyPrice}_s$ is a probability-weighted *average*, and resilience is a property of the tail, not the average: a fleet can hold a flattering $\text{VerifyPrice}_{\text{adj}}$ while its outage-state cost is catastrophic, if the assigned π_s for that state is small enough. The thesis's own treatment of capacity therefore declines expectation forms in the same situation — Delivered Verified Capacity is scenario max-flow/min-cut, with each scenario reported separately, rather than collapsed into a single weighted number (Section 15.9). $\text{VerifyPrice}_{\text{adj}}$ is accordingly a **convenience summary for comparison across fleets, not a quantity any red line or SLO reads on**. The load-bearing publication is the per-state vector $\{\text{VerifyPrice}_s(W)\}$ itself, published state by state alongside the π_s used, so that a skeptical reader can reweight with their own probabilities and see exactly where the summary flatters.

15.10.3 The Decision Rule

Priority ordering matters here, because “resilience” is exactly the kind of word that can be used to excuse blowing through a constitutional target.

Accept a benign-state resilience premium only if **(i)** the worst pre-declared disruption state's VerifyPrice remains inside its per-state bound, **and (ii)** the π_{common} -weighted average falls, **and (iii)** benign-state Physical and Layer 4 VerifyPrice each remain inside the constitutional SLO band.

The gate is the per-state vector, not the average. Clause (i) binds on the **Layer 4 VerifyPrice SLOs**—the constitutional cost band on verifying workload receipts—and not on Physical VerifyPrice, which is an audit-cost measurement with no SLO band to breach; clause (iii) simply requires both metrics to stay inside whatever band applies to each. VerifyPrice_{adj} appears only in clause (ii) and only as a **tie-breaker between fleets that have already passed (i) and (iii)**; a premium that improves the average while worsening the worst disruption state is rejected by clause (i) regardless of what the average says.

The final clause is not negotiable. Resilience is subordinate to the hinge: a fleet so hardened that ordinary users can no longer afford to verify has defeated the purpose of hardening it. Redundancy that breaches the SLO band is not insurance; it is the failure mode wearing insurance as a costume.

This gives the thesis a definite position on “resilience over efficiency,” worth stating plainly because the slogan travels widely and is usually asserted rather than priced:

- A **bounded, measurable** premium that buys a large reduction in tail risk is acceptable, and should be visible in FCR data rather than hidden in operator margin.
- An **unbounded** premium, or one that cannot be shown to reduce disruption-state cost, is not resilience. It is capital misallocation with a security narrative attached.

The FCR fields of Table 15.5, aggregated as the sovereign optionality index of Section 15.8, are precisely the instrument that distinguishes the two cases. Without them, “we are building resilience” is unfalsifiable. With them, it is a number an allocator can check.

Where this sits in the energy-security literature. The trade-off named here — redundancy and optionality priced against the cost of carrying idle capacity — is the central, long-standing question of energy systems reliability economics: the “resource adequacy” debate over how much excess capacity a grid should carry, and who pays for it. That literature (regulatory and academic, from NERC reliability standards through the capacity-market literature) has spent decades converging on the same conclusion this section reaches by a shorter route: resilience that is not metered becomes rent. The thesis does not import the literature’s models, and its FCR/DVC instruments are its own; but the position taken here is not contrarian, and the slogan it prices is the same slogan that literature prices. Analyst commentary used elsewhere in this Part (Doomberg, in Sources) supplies the scenario framing; the regulatory economics supplies the skepticism.

15.11 Operational Patterns: Profiles, Upgrades, and Failure Modes

Layer 0 is not static; hardware evolves, breaks, and gets deprecated. We need patterns for **living with that churn**.

15.11.1 Hardware Profiles and Workload Binding

Each canonical workload W in Layer 4 is associated with one or more **acceptable hardware profiles** ($H_1, H_2, \dots$). Profiles define:

- Minimum performance characteristics (to keep VerifyPrice in target bands).
- Acceptable side-channel leakage budgets.
- Attestation/sampling histories.
- Known caveats (e.g., “avoid profile H3 for workloads with secret inputs; leaks are too strong”).

When PAL compiles a workload or the router assigns work, it can target specific profiles, diversify across profiles, or refuse high-sensitivity workloads to marginal profiles. This is how **Layer 0 informs the market** rather than hiding under it.

15.11.2 Upgrades and Deprecations

Hardware ages; bugs and backdoors are discovered; fabs change hands. Layer 0 needs clear **life-cycle rules**:

Onboarding: New profiles go through a probation period with extra sampling and conservative Work Credit weights.

Deprecation: When a profile is compromised or obsolete, higher layers stop accepting new Work Credits minted from it, risk-flag existing credits, apply collateral haircuts, or impose prospective ineligibility, and publish an incident report at Layer 6.

Migration: Proof factories and corridor operators need technical paths to migrate workloads off deprecated profiles without massive downtime.

In monetary terms, this is the **hardware analog of a bond downgrade**: transparent, describable, and priced, rather than silently swept under the rug.

Non-Discretionary Downgrade Rules

To prevent “ex post discretionary default” critiques, deprecation and haircut decisions follow a **predictable severity framework**, not ad hoc governance.

Process guarantees:

1. **Evidence threshold:** S1 requires independent lab confirmation; S2 requires reproducible demonstration; S3 requires active exploitation observed.

Severity	Definition	Examples
S0 (Watch)	Potential issue; under investigation	Anomalous side-channel readings; unverified third-party report
S1 (Warning)	Confirmed issue with limited impact	RNG bias below 1%; isolated firmware bug
S2 (Critical)	Confirmed issue with systemic impact	Backdoor in >5% of sampled units; key extraction demonstrated
S3 (Emergency)	Active exploitation or catastrophic risk	Widespread key compromise; vendor collusion confirmed

Table 15.7: Severity levels for hardware profile issues.

2. **Multi-party decision:** Severity escalation requires sign-off from ≥ 3 of 5 designated hardware security reviewers.
3. **Appeal process:** Profile sponsors can challenge findings within 14 days.
4. **Sunset, not confiscation:** Even at S3, existing credits remain eligible at 75% of protocol-recognized collateral value. This is a statement about the protocol's own accounting, not a guarantee about market price: no protocol can promise what its asset trades for after a catastrophic hardware compromise, and this document makes no such promise.
5. **Transparency:** All severity determinations are published.

15.11.3 PQ and Cryptographic Agility

Some Layer-0 assumptions are about **cryptography**, not just silicon: signature schemes in ROMs, hash functions in hardware accelerators, and RNG primitives.

Layer 0 insists that hardware and firmware expose enough **configurability** to migrate to post-quantum or new primitives without throwing away entire fabs. Hardware profiles document cryptographic agility (e.g., “can switch hash from X to Y via firmware; signature scheme fixed”).

15.12 Stress Tests for Layer 0

Every claim in Layer 0 should map to a testable stress scenario:

If a stress test fails, Layer 0 **degrades visibly**: profiles are flagged, Work Credits tied to them are risk-flagged and may face prospective ineligibility, and Layer 6 publishes an incident report.

15.13 What Layer 0 Exports to Higher Layers

Higher layers consume a **small set of artifacts and APIs**:

1. **Hardware profiles (HIDs).** Compact identifiers + dossiers describing profile properties,

Claim	Stress Test
Profile H is honest within bounds X	Feed adversarial inputs; measure deviation from spec; decap sample devices
Sampling methodology detects tampering	Red-team: insert known-bad units into supply; measure detection rate
RNG entropy meets threshold	Run standard test suites (NIST, Dieharder); inject synthetic bias; verify detection
Power telemetry is honest	Cross-check metering with grid records; test resilience to meter spoofing
Migration from H1 to H2 works	Simulate H1 deprecation; measure latency and failure rate during migration

Table 15.8: Layer 0 stress tests.

sampling history, and current status.

2. **Attestation receipts.** Machine-level statements binding device → profile, firmware → hash, measurement → nonce/time.
3. **Power and health telemetry.** Streams of power usage, failure rates, and uptime patterns.
4. **Incident and status flags.** Signals like “Profile H3 compromised; do not accept new work.”

Chapter 16

Layer 1: Reachability

All of this math still travels as packets.

Every proof, every private swap, every verified FLOP ultimately crosses a handful of cables and radios that a small number of operators can see and shape. A store of value that survives yield-curve control but dies when a few IXPs collude is not a store of value; it is an overlay.

Comms resistance is the stack's oxygen: the property that air still flows when someone leans on the hose.

Prior art this chapter stands on. The layer's components are deployments, not proposals, and the design borrows deliberately: BIP-324 encrypted P2P transport (shipped in Bitcoin Core 27.0) is the wire-format template for “encrypted and boring”; Tor, I2P, and the pluggable-transport lineage are the first-class transports and obfuscation shims; BOLT12 Offers are the receiver-private rendezvous pattern; and the VerifyReach measurement program of [Section 16.5](#) follows the OONI methodology — volunteer vantage points, canary requests, ASN-stratified public aggregation — that has measured real-world censorship at scale since 2012. The thesis's contribution is not any of these mechanisms; it is binding them into a monetary stack whose falsifiability depends on them, and treating their telemetry as monetary instruments. Full entries in Sources.

16.1 Threat Model: What We Must Survive

The comms threat model has four main faces:

Backbone controls. IP and prefix blocking, BGP blackholing, DNS and SNI filtering, DPI-based resets, QoS throttles on known P2P handshakes.

Exposure and linkage. Static addresses and reusable invoices let adversaries map who is getting paid; mempool and gossip surveillance reveal patterns.

Eclipse and routing capture. A node that only talks to a small set of peers in the same ASN can be effectively isolated.

Platform risk at the edge. App-store takedowns, CDN firewalls, corporate endpoint policies that classify P2P clients as malware.

16.2 Design Rules: Protocol Posture Under Pressure

1. **Transports must be encrypted by default and boring on the wire.** Clear-text, distinctive handshakes invite classification. BIP-324-class v2 encrypted transport is the template.
2. **The stack must be transport-agile.** No single path (TCP/TLS, Tor, I2P, QUIC) should be a single point of failure.
3. **Receiver privacy must be the default.** Payment rendezvous should be “addressless”—BOLT12 Offers, path-blinding, and shielded pools.
4. **Settlement must be refund-safe under squeeze.** Atomic, adaptor-signature flows with clear timeout and refund semantics.
5. **Edge admission must remain open.** No “special” relays or whitelisted entry points.
6. **Comms health is part of the telemetry regime.** If reachability silently degrades, neutrality degrades with it.

16.3 Mechanisms: What Needs to Ship

Encrypted, multi-path P2P transports. Full nodes speak encrypted P2P by default, with Tor and I2P as first-class citizens. Camouflage layers (Noise-style patterns, obfs-class shims) allow relays in high-interdiction ASNs to keep moving traffic.

Receiver-private routing for payments. BOLT12 Offers: receiver-private, reusable invoices. Path-blinding and rendezvous routing ensure intermediaries see only their hop. Federated ecash and shielded pools provide endpoint firebreaks.

Settlement survivability. Adaptor-signature atomic swaps as the default for cross-asset payouts. “Abort & refund” as a first-class, well-signposted action.

Topology and anti-eclipse hardening. Peer-set diversity and rotation; connections spread across geographies and ASNs; gossip protocols favor multiplicity of paths.

16.4 How Layer 1 Anchors the Triad

Communications resilience is not an orthogonal concern; it is the medium in which the triad either lives or suffocates.

For Privacy, encrypted transports and receiver-private endpoints are what keep the network itself from nullifying ledger-level secrecy. Shielded pools and privacy coins are only as private as their membrane to the outside world. If every shielded payout can be traced to a stable IP

graph and a static address, “lawful privacy” collapses into a thin veneer over a rich flow-of-funds analysis. Layer 1’s job is to ensure that settlement paths are as hard to pin down as the flows they carry: private not only in state, but in motion.

For Proofs, communications resilience keeps verification a public act rather than a priestly privilege. A world where only a small set of well-positioned nodes can fetch and check proofs is a world where VerifyPrice has quietly become an internal KPI rather than a public commodity. If any honest machine with modest connectivity can still reach a verifier over at least one path, proof markets and receipt ledgers remain subject to universal scrutiny. When reachability fragments, the “public” in public verification becomes aspirational.

For Compute, Layer 1 ensures that useful-work mining and verified inference do not devolve into “whoever still has a clear line to the router wins.” Proof-of-Useful-Work schemes depend on open admission: a wide, geographically and topologically diverse set of provers and miners competing to satisfy claims. If adversaries can choke ingress to a few ASNs or clouds, PoUW degenerates into a club good. Comms health metrics and open-admission design at the edge keep compute supply neutral and keep the token’s claim (“backed by globally demanded work”) from turning into “backed by whichever datacenter the regulator likes.”

All of this math still travels as packets. Layer 0 keeps the machines honest and powered; Layer 1 keeps them in conversation when it is no longer convenient for them to be. The triad’s monetary ambitions depend on both. Without verifiable machines, we do not know what happened. Without resilient communications, we do not know it in time, or at all.

Residual risks: Encrypted transport and obfuscated routing are ingredients, not complete answers to censorship. Residual threats include traffic analysis (pattern and timing correlation), active probing (protocol fingerprinting despite camouflage), endpoint seizure (physical confiscation of devices), long-term sustained shutdowns (where no transport can help), and legal coercion at the social layer. Layer 1 reduces these surfaces; it does not eliminate them. The honest claim is degradation resistance, not censorship immunity.

16.5 VerifyReach: Communications Telemetry

Layer 1 introduces **VerifyReach** as the communications analogue of VerifyPrice:

- **Reachability metrics:** fraction of vantage points from which key services are reachable.
- **Degradation patterns:** which networks experience blocking or throttling.
- **Transport diversity:** percentage of traffic over each transport class.

Target SLOs:

- Core infrastructure reachable from $\geq 95\%$ of sampled ASNs (uncensored regions).

- $\geq 70\%$ reachability for known-censored regions.
- p95 time-to-first-connection $< 10\text{s}$ uncensored; $< 30\text{s}$ censored.
- No country-level view sees more than 10% persistent reachability degradation without triggering incident handling.

VerifyReach feeds into Layer 6 governance: if reachability collapses in a region, incident response kicks in (alternative transports are promoted, routing is adjusted, and the degradation is visible on public dashboards).

16.5.1 VerifyReach Measurement Specification

Like VerifyPrice, VerifyReach requires a rigorous measurement methodology to prevent gaming and ensure credibility.

Dimension	Minimum Coverage	Rationale
ASNs	≥ 500 distinct ASNs globally	ISP-level blocking
Regions	≥ 50 countries; ≥ 10 per inhabited continent (five-continent convention)	Regional censorship
Censorship regimes	Coverage of known filtering states (CN, IR, RU, etc.)	Worst-case reachability
Network types	Residential, mobile, enterprise, datacenter	Filtering varies by type

Table 16.1: VerifyReach sampling frame requirements.

Sampling Frame:

Vantage Points:

- **Community nodes:** Volunteer-run measurement agents (similar to OONI probes). Incentivized via small WC rewards.
- **Independent labs:** At least 3 organizations (academic, NGO, commercial) run measurement infrastructure. No single operator controls $> 30\%$ of vantage points.
- **Diversity requirement:** Vantage points must span ≥ 20 countries and ≥ 100 ASNs for measurements to be considered valid.

The two threshold levels do different jobs and are not redundant. The sampling *frame* targets ≥ 500 ASNs and ≥ 50 countries; the *validity floor* is the minimum below which measurements are not published at all (≥ 100 ASNs, ≥ 20 countries). Continent arithmetic follows the five-continent convention used throughout this document (≥ 10 per inhabited continent).

Metrics: The ttf and succ targets bind **both**, because they measure different things: ttf percentiles are conditional on successful connection and say nothing about how often connection fails, while $\text{succ}_1/\text{succ}_2$ measure the unconditional success rate and say nothing about how slow the successes were. Neither subsumes the other.

Metric	Definition	Target SLO
$\text{succ}_1(N,R)$	Fraction reaching service N via primary transport within 30s	$\geq 95\%$ / $\geq 70\%$
$\text{succ}_2(N,R)$	Fraction reaching via any transport (incl. fallbacks) within 60s	$\geq 99\%$ / $\geq 85\%$
$\text{ttfc}(N,R)$	Time-to-first-connection (p50, p95), <i>conditional on successful connection</i>	$\text{p95} < 10\text{s} / < 30\text{s}$
failure_class	Taxonomy: DNS, TCP RST, TLS, timeout, active probe	Published per region

Table 16.2: VerifyReach metrics and targets. Targets shown as uncensored / known-censored.

Threat	Mitigation
Spoofed vantage points	Periodic challenges (fetch and sign specific data); anomalous behavior triggers exclusion
Selective treatment	“Canary” requests that should succeed; if canaries fail but targets succeed, point is flagged
Measurement capture	Cross-check results from different operators; divergence > 10 percentage points (measured in absolute percentage points between independently computed aggregates) triggers investigation
Temporal gaming	Continuous measurements (not snapshots); 24-hour rolling averages published

Table 16.3: VerifyReach adversarial robustness measures.

Adversarial Robustness:**Publication:**

- Raw measurement data (anonymized to protect vantage point operators) published daily.
- Aggregated dashboards updated hourly.
- Quarterly reports summarizing regional trends, incidents, and transport effectiveness.
- All measurement code is open-source and reproducible.

Monetary Consequences: If VerifyReach for a region falls below thresholds:

- WC minted by operators in that region may face **regional risk premiums** (higher collateral requirements).
- Corridors primarily serving that region are flagged; users see warnings before transacting.
- Incident response is triggered: alternative transports promoted, routing adjusted.

This makes VerifyReach not just a dashboard metric, but an **input to economic risk pricing**.

Chapter 17

Layer 2: Distribution & Execution

Software distribution is where lofty protocol guarantees meet the boring reality of phones, laptops, and routers that have to run code. In a repression cycle, app stores are throttled, DNS is poisoned, and “safety” policies become chokepoints.

It is far easier to lean on an app store, certificate authority, or CDN than to break a new primitive. If privacy, proofs, and compute are to support an associated base-asset monetary candidate, then the binaries that implement those services must keep moving even when networks are hostile.

Prior art this chapter stands on. The supply-chain invariants below (I1–I5) are not novel claims; they are the TUF/Uptane guarantee set — threshold signing roles, freshness and anti-rollback, reproducibility, compromise response, multi-path delivery — restated for this stack, and the transparency-log mechanism is Certificate Transparency applied to release artifacts, with Reproducible Builds as the verification floor. Layer 2’s contribution is narrower: adapting these to a monetary client whose update path must survive exactly the states in which the asset’s holder-side service is supposed to deliver, and wiring the update channel’s health into VerifyReach telemetry. Entries in Sources under “Software Update Security & Transparency.”

17.1 Threat Model & Objectives

An adversary’s repertoire at this layer:

- App-store removals and policy bans.
- DNS poisoning, SNI and IP blocking, TLS interception.
- BGP hijacks or route blackholing for update servers.
- Certificate revocations; captive portals that substitute their own “secure updates.”
- Targeted developer coercion and “emergency directives.”

Objectives:

- **Reachability:** users can discover and fetch releases despite interference.
- **Authenticity:** users can verify what they fetched without trusting a platform.
- **Safety:** upgrades are atomic, reversible, and survive power loss.

- **Neutrality:** no single jurisdiction or vendor can gate updates.
- **Telemetry without doxxing:** reliability is measurable without surveillance.

17.2 Design Rules

Authenticity before reachability. Releases ship with threshold signatures (3-of-5) from independent maintainers, plus inclusion in an append-only transparency log.

Reproducible builds as default. At least two independent builders reproduce binaries from source; the build process emits deterministic artifact hashes and proofs.

Threshold keys and real revocation. Root keys live offline and are sharded; release keys rotate regularly; revocation is a signed artifact in the transparency log.

Binary transparency anchored to a proof ledger. The update log behaves like certificate transparency for releases, periodically anchored to a neutral proof ledger.

Atomic updates with safe rollback. Clients update into an inactive slot, flip a pointer only after verification, and revert automatically on failure.

Graceful degradation. Offline packages (USB, SD, QR) use the same receipts and keys.

Supply-Chain Invariants (Must Hold)

- I1. Authenticity:** Every release requires threshold signatures (≥ 3 -of-5) from independent maintainers in ≥ 2 jurisdictions, plus transparency log inclusion.
- I2. Freshness:** Clients enforce monotonic version counters. Rollback attacks are detected and rejected.
- I3. Reproducibility:** At least 2 independent builders reproduce each release. Divergence is release-blocking.
- I4. Compromise response:** If a signing key is compromised: revocation within 24 hours, affected releases quarantined, post-mortem within 14 days.
- I5. Multi-path availability:** Releases available via ≥ 3 independent channels. No single channel's failure blocks updates for > 24 hours.

17.3 Reference Distribution Architecture

The architecture is deliberately multi-homed. We assume any single channel can become a chokepoint.

Control plane (discovery and metadata):

- Standard HTTPS endpoints behind diverse DNS providers.
- Tor onion services exposing the same manifests.

- Signed release announcements gossiped over generic protocols (Matrix, Nostr).
- Append-only transparency log with roots checkable by light clients.

Data plane (artifact delivery):

- Traditional CDNs and anycast mirrors in multiple jurisdictions.
- P2P swarms (IPFS-like or BitTorrent-class) using content addressing.
- Community micro-mirrors with signed manifests.
- Offline channels: USB/SD bundles, QR-encoded update chunks, radio relays.

All of these are interchangeable from the client's perspective. Whoever delivers the bytes fastest wins, but no one is trusted beyond the hash: the client verifies the final digest and only then installs.

17.4 Key Management & Release Process

Key management is where many otherwise sophisticated systems quietly re-centralize. A single HSM in a single jurisdiction becomes a soft power lever: compromise the humans or the box, and you own the distribution pipeline.

A capture-resistant release process needs:

- **Public release ceremony.** Changes land on protected branches; CI produces deterministic binaries; independent builders confirm bit-for-bit equality. A quorum of maintainers signs the targets manifest.
- **Emergency freeze and unfreeze.** Any maintainer can propose a freeze, but activation requires a threshold of signatures. Unfreeze requires an explicit, signed artifact explaining the incident.
- **Compromise playbook.** Detect, revoke, rotate, and re-sign last-known-good releases. All steps produce receipts that anyone can audit.
- **Jurisdictional diversity.** Key shards and signers are distributed across legal zones, making it difficult for one government to unilaterally subvert the process.

17.5 Update Economics & Neutrality

Distribution has an economics layer just as proofs and compute do. Instead of pretending bandwidth and storage are free, we let mirrors be market actors: they meter what they serve, emit receipts, and get paid for honest work.

- **Receipts for bytes:** mirrors export metered receipts signed by their nodes; a broker aggregates and pays out via privacy rails with SLA escrow.
- **Open admission:** any party can become a mirror by publishing bandwidth and uptime commitments and staking a small bond.

- **Anti-capture telemetry:** publish house-share, top-N mirror share, geographic/ASN diversity. If a CDN dominates beyond thresholds, route weight decays automatically.

17.6 Fallback Playbooks

Having explicit, rehearsed playbooks keeps censorship events from becoming existential crises: **Soft network blocks (DNS/SNI/IP):** Bias clients toward alternative control-plane channels (DoH/DoT, onions), and toward content-addressed P2P for artifacts.

Hard regional blocks: Rely on offline kits (USB/SD bundles, QR sets) seeded through civic institutions and NGOs.

App-store bans: Publish signed sideload bundles with clear, localized installation instructions.

Key compromise: Hit the freeze button; revoke; rotate; and re-sign last-known-good releases, with all steps recorded in the transparency log.

17.7 Lawful Privacy in Distribution

The same architectural moves that make distribution censorship-resistant also make it more legible to good-faith auditors. Transparent logs, reproducible builds, and signed receipts give regulators something objective to look at:

- Which releases were shipped, and when.
- Which vulnerabilities were patched, and how quickly.
- Whether there were jurisdiction-specific forks or backdoors.

Crucially, auditors can examine this evidence without sitting in the middle of every update flow. The system does not ask to be exempt from scrutiny; it insists that scrutiny be applied at the level of receipts and logs rather than taps and implants.

17.8 Update Telemetry

To make distribution part of the triad's credibility, networks should publish continuously:

- **Update health:** p50/p95 metadata and artifact fetch times; per-path success rates; rollback and signature-failure rates.
- **Transparency evidence:** current log roots; inclusion proofs for each release.
- **Decentralization metrics:** top-N mirror share; geographic and ASN diversity.
- **Key material and incidents:** root fingerprints; release-key roster; revocation lists; incident reports.

Chapter 18

Layer 3: Identity & Claims

Identity in a repression-prone, AI-saturated internet cannot be “a file with your name on it.” It must be capabilities you can prove on demand (age, jurisdiction, uniqueness, solvency, model ownership) without handing over the dossier.

Reputation should be the trail of receipts from prior correct behavior, not a database of personally identifying facts.

18.1 Why Identity Must Decouple from Doxxing

The web’s trust default has flipped: seeing is no longer believing, and platforms strip provenance labels inconsistently. Accounts and government IDs are brittle anchors; they leak power to gatekeepers.

What scales instead is **mechanism over memory**: claims backed by cryptographic evidence any honest party can check cheaply and without permission. This is the same shift the thesis makes for money—verification, not authority, is the arbiter.

18.2 Design Goals

Identity and reputation in this stack satisfy six goals:

1. **Privacy by default, proof by construction.** No standing, globally queryable identity graphs. All predicates are proven as needed, most often with zero-knowledge proofs.
2. **Selective disclosure and unlinkability.** Verifiers learn only the outcome of a predicate, not your name or the rest of your credential bundle.
3. **Cheap, public verification.** Any honest laptop can check evidence quickly, measured the same way we measure proofs with VerifyPrice.
4. **Open admission for issuers and verifiers.** Anyone can issue attestations bound to their own reputation and slashing.
5. **Hardware honesty for machine identity.** Machine claims are anchored in verifiable machines with open designs and sampled supply chains.
6. **Lawful privacy.** Regulators get provable, time-bounded visibility via receipts and viewing

keys, not through custodial chokepoints.

18.2.1 Issuer Pluralism Without Issuer Anarchy

“Open admission for issuers” (goal 4) sounds like chaos: anyone can issue AML credentials, so criminals just issue their own. The resolution is that **issuer sets are policy objects**, not global permissions.

Policy Context	Example Predicate	Issuer Set Definition
AML-screened	“User is not on OFAC/EU sanctions lists”	Threshold of ≥ 2 from {Exchange A, Exchange B, Compliance Provider C, Regulator D}
Age-verified	“User is 18+”	Any licensed identity provider in user’s jurisdiction, or threshold of ≥ 2 from global providers
Accredited investor	“User meets SEC accredited investor definition”	Licensed broker-dealer or accredited investor verification platform
Device integrity	“Device is L0-B or higher”	Hardware attestation from ≥ 2 independent TEE vendors or open-source attestation

Table 18.1: Issuer sets by policy context.

Issuer sets by policy context:

How issuers enter/exit sets:

1. **Stake:** Issuers post the base asset as collateral proportional to the value of credentials they issue. This collateral role is one value-accrual mechanism, not proof of monetary premium. If credentials are later found fraudulent, stake is slashed.
2. **Audit receipts:** Issuers must publish periodic audit receipts proving their verification processes meet policy requirements. Audits are conducted by independent third parties.
3. **Slashing:** If an issuer’s credentials are shown to be systematically false (e.g., issuing “not on sanctions list” to sanctioned entities), their stake is slashed and they are removed from relevant issuer sets.
4. **Jurisdictional diversity:** For high-stakes predicates (AML, accredited investor), issuer sets require representation from ≥ 2 jurisdictions to prevent single-government capture.

How verifiers select policy bundles:

- Verifiers declare which policy contexts they require (e.g., “AML_screened AND Age_verified”).
- The protocol resolves these to issuer sets.
- Presentations must include credentials from issuers in the relevant sets.

- Verifiers can add custom issuer requirements (e.g., “must include credential from Regulator D for EU users”).

Why this isn’t centralization:

- No single issuer is required for any predicate.
- New issuers can join sets by meeting stake and audit requirements.
- Issuer sets are transparent and auditable.
- Users can choose which issuers to use within the set.

This model preserves open admission while ensuring that predicates in regulated contexts are backed by accountable issuers.

18.3 The Identity Kernel

We treat identity as a claim machine rather than a name registry.

Identifiers are cheap, ephemeral, and scoped—pairwise keys or DIDs, not global user IDs.

Credentials are issued by many parties (exchanges, employers, DAOs, regulators, devices). Each is a capability statement, not a dossier.

Presentations are one-off proofs of predicates over credentials: membership, attributes in a range, conjunctions of statements. Typically zero-knowledge.

Receipts are the portable artifacts. Every accepted presentation produces a PIDL receipt binding the claim, proof hash, SLA tier, and timestamps.

This kernel compiles directly into the developer surfaces: PAL SDK to express “prove this predicate,” PIDL receipts to carry results, and privacy corridors to pay without doxxing.

18.4 Human Identity: Personhood and Compliance Without Dossiers

What to prove, privately:

- **Uniqueness / Sybil resistance.** Rate-limited nullifiers, stake-and-slash, or attested-liveness ceremonies.
- **Eligibility.** Age, residency, licensing, or “member-of a KYC-screened set,” proved with set-membership or range proofs.
- **Standing & solvency.** Proofs of reserves/income ranges, revealed under viewing keys during credit underwriting.

How it composes with the stack: A wallet invokes PAL to compile predicates to proofs, pays over privacy rails, and gets a PIDL receipt the counterparty can verify on-chain or in a browser in < 5s p95.

18.4.1 Sybil Resistance Menu

Sybil resistance (preventing one entity from creating unlimited fake identities) is a core challenge. Different mechanisms offer different tradeoffs; the stack supports a menu of approaches rather than mandating one.

Mechanism	How It Works	VerifyPrice	Unlinkability	Best For
Rate-limited nullifiers	Each identity can perform N actions per epoch; nullifiers prevent reuse without linking	Low	High	Spam prevention, voting, airdrops
Stake-based	Identity actions require WC collateral; bad behavior $\rightarrow$ slashing	Low	Medium	High-stakes actions, market making
Social/liveness ceremonies	Periodic video calls, in-person events, or social graph vouching	Medium	Low-Medium	Proof of personhood, high-value credentials
Device-based attestation	Hardware attestation limits identities per device	Low	Low	Machine identity; bounded human use
Economic proof-of-work	Solving computational puzzles to create identity	Medium-High	High	Bot prevention, rate limiting

Table 18.2: Sybil resistance mechanism menu.

Tradeoff guidance:

- **For maximum privacy:** Rate-limited nullifiers + economic PoW. No linkage, but limited action rate.
- **For regulated contexts:** Stake-based + issuer-set credentials. Some linkage to stake address, but accountability.
- **For machine identity:** Device attestation bounded by L0 grade. Accept fingerprinting risk for hardware honesty.
- **For proof of personhood:** Social ceremonies with ZK proofs. Higher coordination cost, but strong Sybil resistance.

Composability: Mechanisms can be combined. A high-stakes action might require:

1. Rate-limited nullifier (prevents spam),
2. Stake (ensures accountability),
3. Issuer credential (meets regulatory requirement).

Each layer adds cost and reduces privacy; the protocol allows verifiers to specify requirements and users to choose compliant paths.

18.5 Machine Identity: Capability, Not Brand

Machines (provers, miners, sensors, model endpoints) need identity, but it should look nothing like a cloud vendor account. Their “identity” is a statement of capability on a particular hardware profile coupled with receipts for past behavior.

When a device participates, it attests to what it is and what it did. Attestations link to open hardware profiles. Those attestations are then wrapped in succinct proofs so any chain can verify them cheaply.

This is **reputation for robots**: SLA-backed capacity that accrues a track record via receipts, not by vendor logo.

18.6 Reputation: Priced Behavior, Not Personal Data

- **Positive reputation** accumulates as fulfilled obligations with proofs: deliveries within SLA, honest inferences, timely settlements, clean audits.
- **Negative reputation** accrues as rejects and slashes: failed audits, invalid proofs, settlement timeouts.
- **Scope and decay** are explicit: reputation is contextual and decays unless renewed.

Because receipts are portable and verifiable, reputation becomes a market primitive. A prover can advertise: “over the last 30 days I delivered PROOF₂₀ within p95 < 1s at cost < \$0.001 with 0.04% failure”—and anyone can check.

18.7 Patterns That Keep Identity Private and Usable

zK-KYC / zK-AML. Users prove membership in a regulated allow-list without revealing which issuer or which entry.

Rate-limited actions. Per-relationship nullifiers cap spam without linking handles.

Credit without doxxing. A lender demands a proof set plus a bond; the loan contract auto-slashes on missed proofs. No one learns the borrower’s legal name unless escrow triggers.

Creator & data provenance. Cameras on open hardware sign origin; receipts embed these proofs in a way that survives platform stripping.

Work credentials for models. Model stewards prove they own specific weights and that inferences came from those weights under defined parameters.

18.8 Failure Modes & Guardrails

Any identity layer that matters will eventually be attacked. What distinguishes a durable architecture is whether failures are observable, bounded, and repairable without appealing to a central authority.

Centralized attesters. If one vendor or agency controls machine attestation or anon-cred issuance, identity becomes permissioned by decree. **Countermeasure:** Pluralism—open hardware profiles, lot sampling, and ZK-wrapped quotes that any verifier can check; multiple issuers whose credentials are interchangeable or combined via threshold logic.

Proof cost creep. If p95 verify time drifts up, only gateways can afford to check receipts. **Countermeasure:** Treat identity workloads as first-class citizens of VerifyPrice; treat regressions as Sev-1 incidents.

Bridge and settlement doxxing. If corridors leak metadata, identity collapses into routing tables. **Countermeasure:** Cryptographic atomicity (adaptor signatures), refund safety, and public settlement telemetry.

Reputation blacklists. Global “bad lists” are honey traps for political pressure. **Countermeasure:** Scoped, revocable policies tied to evidence; actors retain the ability to present counter-receipts.

18.9 How Identity Fits the Modular Stack

Identity and reputation are just another set of workloads in the same loop:

1. **Create / Compute.** A human or machine emits a claim: “I am unique and 18+,” “this device captured this video,” “this miner computed MATMUL_4096 under profile P.”
2. **Prove.** The PAL SDK compiles the predicate into a circuit or useful-work audit, routes it to a prover market, and returns a PIDL receipt.
3. **Settle.** The privacy rails kit executes refund-safe payouts across BTC↔ZEC/XMR corridors or shielded pools.
4. **Verify.** Anyone runs `verify(receipt)` locally, on-chain, or inside another proof.
5. **Telemetry.** VerifyPrice, anonymity-set size, swap success rates, and hardware-profile distributions are public.

The stack stops being “crypto plus KYC” and becomes one fabric where work, truth, and permission all pass through the same narrow gate: proofs anyone can check.

18.10 Identity/Reputation SLOs

For operators and allocators, identity becomes legible through metrics:

- **Predicate VerifyPrice.** The canonical five-field tuple $(p_{50,t}, p_{95,t}, p_{50,c}, p_{95,c}, \text{fail})$ per stan-

dard identity predicate, i.e. median and tail verify time and cost plus failure rate. Target: p95 under a few seconds, sub-cent cost.

- **Unlinkability.** Rate at which presentations can be linked by passive adversaries.
- **Settlement privacy.** Corridor success rates and anonymity-set sizes for identity-linked flows.
- **Hardware honesty.** Share of machine-signed receipts tied to open profiles.
- **Reputation health.** Fraction of decisions based on receipts vs. account flags.

Identity $\neq$ **name**; it is a set of provable capabilities. **Reputation** $\neq$ **biography**; it is a ledger of receipts. Bind both to privacy by default, publish VerifyPrice so anyone can check them cheaply, and anchor machine attestations in verifiable hardware.

18.11 Implementation Sketch

All of the above shows up in three simple patterns that developers wire into their apps:

Human $\rightarrow$ **compliance without names.** A wallet encodes a claim like: $\text{prove}(\text{age} \geq 18) \wedge \text{prove}(\text{residency} \in \text{EU}) \wedge \text{prove}(\text{member_of AML_screened_set})$. It sends this to PAL, which compiles the predicates, sends them to a prover market, and returns a PIDL receipt. The merchant verifies in $<5\text{s}$ p95 without ever seeing the passport.

Machine $\rightarrow$ **capability without brand.** A sensor on hardware profile `open_camera_v1` captures video. The device emits an attested capture, PAL wraps it in a succinct proof, producing a provenance proof and receipt. The camera’s “identity” is simply “a device of profile P that has produced N correct receipts.”

Reputation $\rightarrow$ **receipts, not profiles.** A prover’s handle accumulates receipts. When it advertises itself, it publishes: “Last 30 days: delivered PROOF_2²⁰ under SLA Gold with p95 verify-time 0.8s, 99.9% success, zero slashes.” Anyone can pull the underlying receipts and run the reference verifier.

18.12 Part III Summary: Minimum Viable Stack

Before proceeding to the economic layers (4–5), we consolidate the **minimum viable requirements**:

Telemetry Reproducibility Principle

Telemetry must be reproducible from raw public artifacts by multiple independent parties, or it is not part of the trust base.

Layer	Minimum Viable Requirement	Falsification Test
L0	At least one L0-B grade profile; sampling for L0-A	If proving concentrates on L0-A hardware and lot-attestation coverage stays below the L0-B threshold ($\geq 80\%$ of production lots) $\rightarrow$ thesis weakened
L1	Encrypted P2P + ≥ 2 fallback transports; VerifyReach with ≥ 3 operators	If $> 20\%$ ASNs unreachable for > 7 days $\rightarrow$ thesis fails
L2	Threshold-signed releases + transparency log + ≥ 2 builders	If single key compromise can push malicious updates $\rightarrow$ thesis fails
L3	≥ 3 identity predicates with $p_{95,t}$ of the canonical tuple $< 5s$; issuer pluralism	If identity requires centralized issuer $\rightarrow$ thesis weakened

Table 18.3: Minimum viable requirements for Layers 0–3.

Specifically:

1. **Raw data availability:** All inputs to telemetry metrics are published or archived.
2. **Open methodology:** Aggregation formulas and statistical procedures are documented.
3. **Multi-party verification:** At least 3 independent parties publish telemetry; divergence > 10 percentage points (absolute, between independently computed aggregates) triggers investigation.

Four layers now stand underneath the economic ones: matter that can be audited, packets that survive pressure, code that ships without a gatekeeper, and identity that travels without a name. Each was specified so that its failure is a measurement, not a debate. Part IV asks the only question left for a stack this deep: what work is worth settling on it.

Part IV
Truth, Work, and Settlement: Layers
4–5

Part III started at the basement: Layer 0 as verifiable machines and energy. Layers 1–3 keep packets moving, code running, and actors named without being doxxed.

Part IV is where the stack becomes **explicitly economic**. Layers 4 and 5 are the point where:

- raw compute becomes **truth** (via proofs and verification asymmetry), and
- truth becomes **money** (via private, non-custodial settlement).

This is the layer pair that turns the abstract triad (Privacy, Proofs, Compute) into something you can actually buy and sell: **proof commodities, verified FLOPs, and privacy-preserving flows**.

Chapter 19

From Infrastructure to Economics

Layers 0–3 answer four questions:

1. *Can we trust the machine?* (Layer 0)
2. *Can we reach it under censorship?* (Layer 1)
3. *Can we get code onto it and keep it updated?* (Layer 2)
4. *Can actors prove “who” they are without doxxing themselves?* (Layer 3)

Layers 4 and 5 take those answers as given and ask two more:

5. *Can we turn machine work into claims that anyone can verify cheaply?* (Layer 4)
6. *Can we settle those claims as value flows without chokepoints?* (Layer 5)

If Layers 0–3 are the **nervous system and limbs**, Layers 4–5 are the **cortex and circulatory system**: they decide what counts as action and how it is remembered in economic form.

Infrastructure $\neq$ Money

Layers 4–5 create markets for proofs, verified compute, and private settlement. They still do not create money unless Part II’s value-capture conditions hold. A working market for verified FLOPs is necessary for the monetary thesis but not sufficient; the base token must capture the economics of that market through fees, burns, and collateral as specified in the [Value Capture Lemma](#).

Chapter 20

Layer 4: Truth & Work

Layer 4 is the **Prove/Verify spine** of the stack.

Its job is to take arbitrary claims—“this media object came from camera C at time T,” “this inference was computed by model M on input X,” “this corridor swap executed atomically”—and turn them into **portable proofs** that:

- anyone can check with **cheap, public verification**, and
- can be **standardized as commodities** (canonical workloads, SLAs, prices).

20.1 What Layer 4 Is (and Isn't)

Layer 4 is:

- The layer of **circuits, proofs, and verification economics**.
- The place where **verification asymmetry** is engineered and measured.
- The home of **Proofs-as-a-Library (PAL)**, proof factories, and **canonical workload registries**.

Layer 4 is not:

- A specific proof system (SNARK vs. STARK vs. something else); it assumes **multi-ZK**.
- A single PoUW design; it supports several patterns as long as VERIFYPRICE and anti-capture constraints are met.
- A truth oracle. Proofs do not prove truth; they prove specific claims under stated assumptions—origin, custody, computation, policy compliance, and settlement finality. The thesis is not that cryptography tells us what is true, but that it reduces the surface area over which institutions must be trusted.

20.2 Verification Asymmetry Revisited

Recall the definition from Part I. For a workload W :

- $p(W)$ = cost (time, energy, hardware) to **produce** a result + proof.
- $v(W)$ = cost to **verify** that result + proof.
- $r(W) = \frac{v(W)}{p(W)}$ = **verification asymmetry**.

Both $v(W)$ and $p(W)$ are aggregated into the same numeraire — USD via the per-run price

model of Section 20.6 ($c_{\text{run}} = (e_{\text{run}}/3.6 \times 10^6) p_{\text{energy}} + t_{\text{run}} p_{\text{opp}}$) — so the ratio is dimensionless and both halves carry the same measurement contract: v under the adversarial harness, p under the reference production model (Appendix A). Read the denominator before celebrating the ratio: $r(W)$ improves in two ways, verification getting cheaper or production getting dearer, and a passing $r(W)$ achieved by expensive production is a cost problem wearing an asymmetry credential. The ratio is therefore published with both components; a falling $r(W)$ driven by a rising $p(W)$ is flagged as such.

Layer 4's goal is simple:

Make $r(W) \ll 1$ for the workloads that matter, and keep it that way in production.

This is the touchstone of the gold analogy, now carrying a price: expensive to smelt, cheap to assay. $r(W)$ is the assay, priced.

Why this matters:

- If $v(W)$ is small and stable, **anyone** (including small nodes) can check proofs.
- That makes proofs and verified FLOPs **commodities**: units of work that any counterparty can accept without trusting a platform.

Layer 4 introduces:

$\text{VerifyPrice}(W) = \{\underbrace{\text{p50 time, p95 time, p50 cost, p95 cost, failure rate}}_{\text{canonical five-field tuple (Appendix A)}}, \text{hardware profile mix}\}$

where the first five fields are the canonical tuple of Appendix A and the hardware profile mix is Layer 4's reporting context: every published series names the verifier classes it aggregates over, and the mix travels with the tuple rather than replacing any field.

20.3 Canonical Workloads and Proof Types

To avoid an unbounded zoo of bespoke proofs, Layer 4 maintains **canonical workloads**:

- **MatMul**($n, k, m; \varepsilon_r$) – matrix multiplication at specified dimensions and relative error bound.
- **Inference**($M, X; \text{policy}$) – run model M on input X under constraints.
- **Provenance**(C, chain) – provenance chain for content type C .
- **Settlement**(S, policy) – settlement and refund logic for a corridor.

For each canonical workload, the stack defines:

- **Circuits / arithmetization**: how the workload is represented for proving systems.
- **Proof schemas**: which proof systems are supported.
- **SLO tiers**: latency and reliability classes (Bronze, Silver, Gold).

Canonical workloads matter economically:

- They become the **SKUs** for proof and compute markets.

- Work Credits are minted against units of these workloads.
- VERIFYPRICE and $r(W)$ are tracked per workload and tier.

The precedent for standardized work. The requirement that demand standardize into canonical units before it can become an economic object has a long pedigree, and it is worth naming because it is easy to mistake for an engineering preference. The standardization thesis in the print literature (Eisenstein 1979; Sources) documents the sequence: typographic uniformity—every letter identical to every other letter of its kind, modular, movable, recombinable—was daily practice at scale for roughly three centuries before the machine shop embodied the same template in screws, gauges, and interchangeable parts. The claim is causal and well-documented as mechanism: standardization of the cognitive module preceded and enabled standardization of the physical one. The stronger claim—that industrialization was possible *only* because of movable type—is not adopted here. The analogy to this layer is exact at the level of function: canonical workloads and PIDL receipts are what make two providers’ units of work identical, comparable, and combinable; diffuse demand for “compute” becomes an economic unit only when the units are interchangeable. The registry is the type case; the receipt is the printed page.

20.3.1 Canonical Workload Definition Template

For a workload to become a tradable SKU (and eligible for Work Credit issuance), it must be defined precisely enough that two providers’ units are interchangeable—and that is the whole reason this template exists. Free-form compute cannot be a unit of account. If one vendor’s “matmul” means fixed dimensions under a relative error bound and another’s means whatever their kernel happens to produce, no counterparty can compare the two, no market can price them, and no registry can mint against them. The template fixes the fields that turn work into a unit: what is claimed, what is public, what is private, what gets checked, what it costs to check, and where it may run.

One field deserves to be stated in the open rather than left inside a parenthesis, because a natural first draft of this template gets it wrong. The error bound is *relative and scale-normalized*, never absolute. An absolute bound would fail honest provers: standard BLAS kernels reorder floating-point arithmetic, and the discrepancy that reordering introduces grows with the magnitude of the operands, so a correct product computed by a different honest kernel would still fail the test. The relative bound forgives honest implementations and still fails results that are actually wrong. Strict with dishonesty, tolerant of implementation—that is the register the entire template is written in.

Canonical Workload Definition Template

WorkloadID: Unique identifier (e.g., MATMUL_4096_FP32, INFER_LM_70B_256TOK)

Statement: What is being proved (natural language + formal predicate)

- *Example:* “The matrix product $C = A \times B$ was computed correctly, where $A, B \in \mathbb{R}^{4096 \times 4096}$, and $\|C - A \cdot B\|_\infty \leq \varepsilon_r \|A\|_\infty \|B\|_\infty$ ” — the *relative, scale-normalized* bound of Section 20.3; an absolute bound is not the template’s form, because honest provers running standard BLAS kernels fail it.

Public Inputs:

- Commitment to A, B (hash or Merkle root)
- Claimed result commitment (hash of C)
- Error bound ε
- Timestamp range

Private Inputs (Witness):

- Full matrices A, B, C
- Intermediate computation trace (if required by proof system)

What Is Verified:

- Correctness (computation matches claim)
- Bounds (result within specified limits)
- Freshness (timestamp within allowed window)
- Liveness (proof generated within epoch, not precomputed)

Verifier Complexity Class:

- Time: two regimes, and the registry entry must say which it prices. *Succinct-proof* workloads verify in $O(\log n)$ after the committed output is in hand; *Freivalds-style* audit workloads verify in $O(n^2)$ but avoid the $O(n^3)$ recompute (this is the MATMUL class of Section 20.3). The template’s complexity field names the regime; a workload carrying both regimes publishes both.
- Memory: specified peak (e.g., $\leq 512\text{MB}$ for Laptop-Class)
- Proof size: max allowed (e.g., $\leq 1\text{MB}$)

Policy Hooks: (machine-verifiable predicates)

- Hardware profile requirements (e.g., L0-B or higher)
- Prover stake requirements
- Membership/non-membership predicates (allowlist proofs, not graph inspection)

Allowed Hardware Profiles:

- L0-A: Yes (with issuance weight $0.9\times$)
- L0-B: Yes ($1.0\times$)
- L0-C/D: Yes ($1.1\times$)

Tier	Latency Target	Redundancy	Fee Multiplier
Bronze	$p95 \leq 60s$	$1 \times$ verification	$1.0 \times$
Silver	$p95 \leq 10s$	$2 \times$ verification	$1.5 \times$
Gold	$p95 \leq 2s$	$3 \times$ verification + audit	$2.5 \times$

Table 20.1: SLA tiers for canonical workloads. The latency target is the *prove-side* time to produce the attested result (generation plus proof), not verify-side time: verification SLOs are constitutional and set separately in [Section 20.7](#), and no prove-side tier can relax them. The redundancy column is likewise prove-side (independent re-executions for attestation); verify-side cost is single-checking by construction, which is the asymmetry the chapter is about.

SLA Tiers:

Field	Value
WorkloadID	MATMUL_4096_FP32
Statement	Matrix multiply $C = A \times B$, dimensions 4096×4096 , FP32, relative error $\ C - A \cdot B\ _\infty \leq \varepsilon_r \ A\ _\infty \ B\ _\infty$, $\varepsilon_r = 10^{-3}$
Public Inputs	Hash(A), Hash(B), Hash(C), ε_r , reference kernel ID, timestamp
Private Inputs	A, B, C , intermediate products
Verified	Correctness, Bounds, Freshness
Verifier Complexity	$O(\log n)$, $\leq 256MB$, $\leq 500KB$ proof (receipt; the output artifact C is $\sim 64MB$ and is delivered separately from the receipt)
L0 Requirement	L0-A minimum; L0-B+ for Gold tier
VerifyPrice Target	$t_{95} \leq 2s$ (Laptop-Class)

Table 20.2: Example canonical workload: MATMUL_4096_FP32. The acceptance bound is *relative*, scaled by $\|A\|_\infty \|B\|_\infty$: FP32 accumulation over $k = 4096$ terms typically incurs $\sim 10^{-4}$ – 10^{-3} relative rounding error with standard BLAS kernels, so an absolute bound of 10^{-5} would fail honest provers. The reference kernel ID pins the arithmetic order so heterogeneous provers produce comparable transcripts; see the determinism note in [Section 20.3](#).

Example: MATMUL_4096_FP32 **Why this matters:** Without this template, “canonical workload” risks becoming “whatever the prover says it is.” With it, workloads are **precisely specified** (anyone can implement a conforming prover/verifier), **auditable** (claims can be checked against the template), and **tradable** (markets can price and exchange standardized units).

20.4 PoUW Design Patterns

Layer 4 doesn’t pick a single consensus recipe. It supports several **proof-of-useful-work patterns**, as long as they satisfy:

- **Open admission:** commodity participants can join.
- **Unpredictable leader election:** no one can cheaply bias the lottery.
- **Useful work binding:** you can’t precompute offline or reuse stale work.

- **Proof quality & anti-spam:** junk proofs can't flood the system without penalty.

Pattern 1: Hash-gated useful work

- Miners perform a cheap hash race; crossing a threshold gives short-lived eligibility.
- To produce a valid block, the miner attaches a PoUW artifact seeded from header randomness.

Pattern 2: Proof-first selection

- Provers race to produce useful-work proofs and post them to a mempool.
 - A lightweight mechanism selects which proofs get included and rewarded.
- Both patterns are compatible with MatMul-PoUW, Inference-PoUW, and hybrid schemes.

PoUW Security Properties (Must Hold)

- P1. Precomputation Resistance:** Proofs seeded by unpredictable randomness; valid only for $N \leq 10$ blocks after seed.
- P2. Grinding Resistance:** Attacker with $X\%$ hashrate gains at most $(1 + \delta) \cdot X\%$ rewards, where $\delta \leq 0.1$.
- P3. Network Advantage Bound:** p95 propagation delay $\leq 2s$ for proof announcements.
- P4. Spam Resistance:** Deposit $\geq 10\times$ expected verification cost; 100% slash for invalid proofs.
- P5. Cartel Detection:** Top-1 share $< 20\%$; Top-5 share $< 50\%$; entry latency ≤ 7 days.
- P6. Useful Work Binding:** $\geq 10\%$ of proofs independently re-verified by random auditors.

Component	Maturity	Status
MatMul-PoUW verification asymmetry	Experimental	Conjecture (arXiv); not yet deployed
Inference PoUW (Proof-of-Logits)	Pilot	Tested in limited networks
Full ZKML for large models	Research	Not yet practical at production scale
Hash-gated useful work (Pattern 1)	Experimental	Devnet demonstrations
Proof-first selection (Pattern 2)	Experimental	Early implementations

Table 20.3: PoUW component research maturity.

Research Maturity Labels: Implication for service claims: The “AI Money” label is an analytical lens on verified-compute demand, not an instrument classification. As components mature from Experimental to Production, inference Work Credits can receive stronger assurance grades and smaller service haircuts. Their issuance remains separate from the base-asset monetary constitution (see Tier taxonomy in [Section 22.1](#)).

Verification Type	Service-Claim Eligibility
Full cryptographic proof (Tier A)	Full eligibility for high-assurance Work Credit issuance
Probabilistic audit with explicit error bounds (Tier B)	Discounted or limited eligibility
TEE/vendor attestation only (Tier C)	Service market only; not eligible for high-assurance capacity claims
Unverified output	No eligibility

Table 20.4: Verification type and typed service-claim eligibility.

20.5 Proof Factories and PaL

Most developers don't want to think about circuits; they want to say:

"Prove that this computation happened, then get me a receipt and pay whoever did the proving."

Layer 4 provides this via:

- **Proof factories:** infrastructure clusters specialized in generating proofs.
- **Proofs-as-a-Library (PaL):** an SDK that compiles high-level claims into proofs.

PaL exposes interfaces like:

```
prove_compute(f, inputs, policy)
prove_provenance(asset_id, lineage)
prove_settlement(tx, corridor_policy)
```

Under the hood, PaL:

1. Maps the request to a canonical workload W .
2. Selects suitable proof systems and hardware profiles.
3. Submits the job via neutral routers to proof factories.
4. Returns a PIDL receipt plus a proof artifact.

20.6 VerifyPrice Observatory

Verification asymmetry is a design goal; VERIFYPRICE turns it into a dashboard.

The **VerifyPrice observatory** continuously measures:

- p50/p95 verify times,
- estimated energy per verification,
- failure rates and mismatch rates,
- diversity metrics (how many independent verifiers are actually checking).

Class	Hardware Spec	Use Case
Laptop-Class	4-core CPU, 16GB RAM, SSD, no GPU	Baseline for “anyone can verify”
Mobile-Class	ARM SoC, 8GB RAM, flash	Edge verification, IoT
Server-Class	32-core CPU, 128GB RAM, optional GPU	High-throughput

Table 20.5: Reference verifier classes for VerifyPrice measurement.

Reference Verifier Classes:**20.6.1 VerifyPrice Measurement Specification**

VerifyPrice is a service-assurance KPI and one necessary input to the base asset’s conditional monetary test. It determines whether Work Credits retain a verifiable service proposition; it does not make them monetary. That requires a rigorous measurement harness, not just a dashboard slogan.

The harness exists because the observatory is not a disinterested instrument. It is a publisher with interests, run by operators who can be pressured, and the numbers it publishes are exactly the kind a protocol under strain would prefer not to publish. Assume the observatory is fallible, and further assume it is *adversarially* fallible—inclined toward flattering measurements when the truth is inconvenient. The specification below is not written to make the observatory honest. It is written so that when the observatory lies, anyone can catch it: the samplers are named, the harness is open, the raw data is archived, and independent operators run the same tests on the same proofs. A lie has to survive all of that. The spec does not produce trust in the measurement; it produces detection of the absence of it.

Cost Vector Definition: For each workload W and verifier class V , VerifyPrice is a 5-tuple:

$$\text{VerifyPrice}(W, V) = (t_{50}, t_{95}, e, m, f)$$

Where:

- t_{50}, t_{95} : median and 95th-percentile verification **time** (seconds)
- e : **energy** per verification (Joules, measured via hardware counters or watt-meter)
- m : peak **memory** (MB)
- f : **failure rate** (timeouts, invalid proof rejections, crashes)

USD cost is derived *per run*, then percentiled: $c_{\text{run}} = \frac{e_{\text{run}}}{3.6 \times 10^6} \times p_{\text{energy}} + t_{\text{run}} \times p_{\text{opp}}$, where e_{run} is the run’s measured energy in Joules, p_{energy} is the reference energy price (USD 0.10/kWh, so the first term divides by 3.6×10^6 J/kWh), and p_{opp} is opportunity cost (USD 0.001/s), both published quarterly. The cost percentiles of the canonical tuple — $p_{50,c}$ and $p_{95,c}$ — are the percentiles of c_{run} across the measurement window, never recomputed from percentile inputs; in particular, deriving a cost figure from t_{95} alone is not the construction.

Requirement	Rationale
≥ 2 independent implementations per workload tier	Prevents single-implementation bugs from corrupting measurements
Open-source, reproducible builds	Anyone can audit and rebuild
Deterministic output	Same proof $\rightarrow$ same result, always
Versioned and tagged	Measurements tied to specific verifier version

Table 20.6: Verifier implementation requirements.

Verifier Implementations:**Sampling Methodology:**

- **Random selection:** Proofs selected uniformly at random from recent submissions (not cherry-picked).
- **Stratified by region/profile:** Measurements cover ≥ 10 regions and ≥ 3 hardware profiles per workload.
- **Adversarial corpus:** 10% of test proofs are malformed or worst-case (max witness size, pathological inputs) to measure failure handling.
- **Continuous measurement:** Not periodic snapshots; rolling 24-hour windows published hourly.

Condition	How Tested
Network latency	200ms RTT injected for proof fetch
Packet loss	5% random packet loss during verification
Worst-case proof size	Max allowed witness for workload class
Malformed proofs	10% of test corpus intentionally invalid
Resource exhaustion	Verification under 80% memory pressure

Table 20.7: Adversarial conditions for VerifyPrice testing.

Adversarial Conditions:**Reproducibility Standard:**

- **Benchmark harness:** Open-source, versioned, deterministic. Anyone can run the same tests.
- **Signed results:** Each measurement batch is signed by ≥ 2 independent measurement operators.
- **Divergence alerts:** If independent operators diverge by $>5\%$, investigation triggered.
- **Archived raw data:** All proof samples and timing logs archived for 1 year.

Why this matters: Without this spec, “VerifyPrice observatory” is an oracle claim. With it, VerifyPrice becomes **reproducible consensus**—any skeptic can run the harness, check the measurements, and falsify the dashboard if it’s wrong.

Workload	Metric	Target (Laptop)	Sev-1 Threshold
PROOF_2 ²⁰	t_{95}	$\leq 5s$	$> 10s$ for 7 days
MATMUL_4096	t_{95}	$\leq 2s$	$> 5s$ for 7 days
INFER_LM_7B	t_{95}	$\leq 30s$	$> 60s$ for 7 days
INFER_LM_70B	t_{95}	<i>aspirational; not currently achievable</i>	n/a until ZKML matures
All	failure rate	$\leq 0.1\%$	$> 1\%$ for 7 days

Table 20.8: Constitutional VerifyPrice targets. The 70B-inference row is deliberately marked aspirational: current ZKML systems prove quantized transformer inference at small scale with substantial overhead, and no production system verifies a 70B-parameter forward pass on laptop-class hardware within seconds. Placing it in the constitutional set as though it were deliverable would trip Red Line 1 on publication day. It remains in the registry as the commercial destination; it is not a constitutional SLO until the maturity gate of [Section 20.10](#) lifts for its component.

Physical VerifyPrice SLOs (Constitutional):

20.7 VerifyPrice in Practice

At the North-Star level, VerifyPrice answers a simple allocator question:

“If I hold this asset for a cycle, does one unit still buy at least as much verification as it used to?”

We care about VerifyPrice in three dimensions, designed to avoid circular reasoning:

20.7.1 The Constitutional SLO Band

These are **non-negotiable** targets, measured in real resources on reference hardware. They are **exogenous to token price**. Whether the token rallies or dumps, these targets must hold. If verification takes 30s on a laptop, the “anyone can verify” promise is broken. The table below is the one referenced as [Section 20.7](#)’s SLO band; the measurement specification that produces its inputs is [Section 20.6](#).

What “Exogenous” Does and Does Not Cover

Exogeneity to token price is the property we need here, and it is the one being claimed: the target must not be redefined because the market moved.

It is not exogeneity to everything. Measuring in real resources on reference hardware makes these SLOs dependent on two conditions a state can adjust—whether unprivileged reference hardware remains obtainable, and what power costs at the edge and at the prover. Both are policy surfaces ([Section 5.5](#)).

The consequence is that Red Line 1, the thesis's hinge, is externally triggerable. This is monitored upstream via sovereign optionality (Section 15.8) and Red Line 13 rather than left to be discovered at the hinge. Section 15.10.1 separates the verification-side exposure (hardware obtainability) from the proving-side exposure (bulk energy), which behave differently and should not be conflated.

There is also a reason the exogeneity requirement runs deeper than engineering preference. If the framing of Section 2.0.1 is right that a medium restructures only those who *habituate* to it, then verification must be cheap enough to become a habit before it can become anything else. If its cost tracks the token's price, habit never forms, and neither does whatever the habit would have built. The SLO is not a performance target; it is the precondition of the medium having any effect at all.

20.7.2 Protocol Affordability SLOs (Operational)

These measure whether verification remains affordable as a fraction of typical transaction costs:

$$\text{AffordabilityRatio}(W) = \frac{\text{VerifyCost}(W)}{\text{MedianFee}(W)}$$

- **Target:** $\text{AffordabilityRatio} \leq 5\%$ for all workload classes.
- **Settlement:** Verification cost $\leq 1\%$ of median transaction value.

Read the denominator before celebrating the ratio. The ratio improves in two ways: verification gets cheaper, or the fee gets bigger. A protocol that raises $\text{MedianFee}(W)$ passes this SLO while becoming less affordable to use, so the ratio must never be published without its components, and a falling ratio driven by a rising denominator is a deterioration wearing an improvement's clothes. The settlement variant has the mirrored exposure: if median transaction value is measured in token terms, appreciation mechanically lowers the ratio and re-imports token price through the back door — the exact circularity the physical SLOs exist to exclude. Three guards hold the family honest, each of a different kind. The first is component disclosure: the fee-based ratio is always published alongside the fee level itself, so a denominator-driven improvement is visible on inspection. The second is a unit convention: the settlement ratio is computed on USD-referenced transaction value, never token-referenced. The third is a monotonicity rule: the affordability SLO is credited only when the ratio falls while the fee level is flat or falling, and improvements achieved on a rising fee are reported as neutral, not as passes. The affordability family is operational, not constitutional, and these guards are why it stays that way honestly.

20.7.3 Token-Quoted VerifyPrice (Market Signal, Not Target)

Token-quoted VerifyPrice is a **useful market signal**, but not a constitutional target. For a scalar reading, fix a single reference workload W^* and take the cost-percentile field only, with the fee schedule expressed as a fixed token quantity:

$$\text{VerifyPower}(\text{token}; W^*) = \frac{q_{\text{fee}}(W^*) \cdot P_X}{p_{50,c}(W^*)}$$

where $q_{\text{fee}}(W^*)$ is the fee schedule for W^* denominated in units of the native asset X , and P_X is the asset's USD reference price at the observation — so the signal carries the token price explicitly on the right-hand side, which is what makes it a *market* signal rather than a disguised operational one. A vector has no reciprocal; the scalar must be constructed by naming the workload and the field. Which workload is the reference is a publication choice, not a free parameter — it is declared, pinned, and changed only with notice, so the signal cannot be dressed by re-picking W^* . Note the family relationship: this quantity is the token-side image of the AffordabilityRatio, and the two must be published together so neither is read alone; the affordability ratio is the operational SLO, this is its purchasing-power translation.

Token price is endogenous. Making this a *target* creates circular reasoning. The protocol commits to **physical SLOs and affordability ratios**. Token-quoted metrics are dashboards for market participants, not governance constraints.

Level	What It Measures	Who Enforces	Consequence of Breach
Physical (Constitutional)	Real-world verification cost	Protocol governance	Sev-1; remediation required
Affordability (Operational)	Verification as % of fees	Fee policy	Fee schedule review
Token-quoted (Market)	Purchasing power signal	Market participants	Informational only

SLO Hierarchy:

20.8 Physical VerifyPrice: Auditing the Infrastructure Behind a Workload

Note on terminology: Section 20.7 uses “Physical VerifyPrice SLOs” to mean the real-resource cost of verifying a *proof or compute claim* on reference hardware, as distinct from token-quoted metrics. This section introduces a related but distinct metric: the cost of verifying the *physical infrastructure claims* (energy, hardware, jurisdiction) behind that proof or compute claim, i.e., the Facility Capacity Receipts introduced in Section 15.7.

VerifyPrice tells us whether a proof or compute claim is cheap to check. **Physical VerifyPrice** tells us whether the physical claims behind that proof are cheap to audit. A Work

Credit cannot be high-assurance if verifying the proof takes seconds but verifying the facility, hardware profile, energy receipt, and jurisdictional risk takes months of trust-the-vendor diligence.

Concretely, Physical VerifyPrice asks:

- Can an allocator verify the facility’s claimed energy profile?
- Can an auditor verify outage and curtailment history?
- Can the network verify that the hardware profile was not silently substituted?
- Can a user verify that compute was not routed through a sanctioned or coerced jurisdiction?

Physical VerifyPrice

The time, cost, and confidence required to verify the physical infrastructure claims (energy, hardware, jurisdiction) behind a unit of verified work.

For each facility profile, the stack should publish p50/p95 time and cost to verify FCR claims, sampling frequency, third-party challenge success rates, and unresolved discrepancy counts. If Physical VerifyPrice rises above threshold, credits from that facility receive risk haircuts or issuance caps (Section 23.6.4). This metric is formalized alongside the VerifyPrice model in [Appendix A](#).

20.9 Layer-4 Stress Tests

Layer 4 passes its SoV audition if it survives several adversarial scenarios:

Circuit bloat. Can we detect when proofs become too expensive to verify? Is there a migration path to leaner circuits?

Prover cartel. Can neutral routers and SLA/slashing mechanisms prevent a small set of proof factories from monopolizing high-value workloads?

Proof system break / new attack. Can we deprecate a proof system, rotate to alternatives, and quarantine affected Work Credits?

If Layer 4 remains **cheap to verify, open to participate, and instrumented enough to handle change**, then Proofs and Compute deservedly move closer to “monetary primitive” rather than “platform feature.”

20.10 Minimum Viable Layer-4 Economy

Canonical Workload Starter Set:

Maturity gate. No workload may issue Tier A (full-weight, pristine-collateral) Work Credits while the maturity label of its verification component (Table 20.3) is Experimental or Pilot.

WorkloadID	Description	VerifyPrice Target	WC Tier
PROOF_2^20	Generic ZK proof	$t_{95} \leq 5s$	A (full)
MATMUL_4096_FP32	Matrix multiply 4096×4096	$t_{95} \leq 2s$	A (full)
INFER_LM_7B_512TOK	7B param inference	$t_{95} \leq 10s$	B (discounted)
PROVENANCE_MEDIA	Media provenance chain	$t_{95} \leq 3s$	A (full)
SETTLE_ATOMIC	Atomic swap settlement	$t_{95} \leq 5s$	A (full)

Table 20.9: Canonical workload starter set. Two target levels are distinct by design and must not be read as one: the constitutional floor (Section 20.7 SLO table) is the minimum a network must hold to remain constitutional, while this starter-set target is the commercial bar the registry holds itself to for the named workload variant (e.g. INFER_LM_7B_512TOK at $t_{95} \leq 10s$ against the 30s constitutional floor for the INFER_LM_7B class). WC tiers shown are the *ceiling* each workload can reach once its verification component is Production-mature; current effective tier is the lower of the ceiling and the maturity-gated tier (Section 20.10 tier rules).

Until a component reaches Production, its workloads issue at Tier B or below regardless of the mathematical tier their proof type would support. As of this writing that gate is binding on every PoUW-based row in the starter set: MatMul-PoUW verification asymmetry is a conjecture, inference PoUW is pilot-stage, and pricing pristine collateral over either would define a AAA rating on an instrument whose proof system has not survived production. For starter-set rows whose verification component does not appear in Table 20.3 at all (PROOF_2^20, PROVENANCE_MEDIA, SETTLE_ATOMIC), the gate reads on the maturity of the underlying proof system family — generic succinct proofs, provenance attestation, atomic-swap verification — each of which the maturity table tracks at the family level; a family not yet tracked is treated as Pilot until it is. The gate lifts only when the maturity table lifts, and the maturity table is moved by deployment evidence, not by argument.

Tier	Verification Type	WC Weight	Collateral
Tier A	Full cryptographic proof (ZK-SNARK/STARK) <i>and</i> component at Production maturity	1.0×	Pristine
Tier B	Probabilistic verification (audited transcripts)	0.6×	Standard
Tier C	Attestation-backed (TEE + sampling)	0.3×	Discounted

Table 20.10: Work Credit issuance tiers.

Tier Rules for Work Credit Issuance:

Fee + Burn + Slashing Logic:

- 70% of fee to prover (reward)
- 20% burned (supply reduction)
- 10% to protocol treasury (security budget)

The 20% burn share here is a deliberately conservative variant of the 30–50% reference band in Chapter 7: this worked example also funds a protocol treasury, so the three shares are

set to sum to 100%. The reference band remains the design target; a 20% burn is the harder case for the value-capture claim, not the easier one.

Fee formula (reference):

$$F(W) = \text{BaseFee}(W) \times (1 + \text{CongestionMultiplier}) \times \text{SLATierMultiplier}$$

Chapter 21

Layer 5: Value & Settlement

Layer 4 turns work into truth. Layer 5 turns truth into **money movement**.

It is the layer where balances change, salaries get paid, cross-asset swaps execute, and collateral gets posted or released.

The constraints are harsh:

Repression-resilience. Default assumption: yield-curve control, capital controls, KYC choke-points.

Privacy & lawfulness. Users need *privacy by default* and *auditability by consent*.

Non-custodial safety. Settlement protocols must be self-custodial and refund-safe.

Privacy Rails and Lawful Obligations

Privacy rails are designed for lawful confidentiality and user-controlled disclosure; they are not framed as tools for evading legal obligations. Some regulatory demands may be incompatible with privacy-preserving systems; in those cases, the protocol acknowledges the tension rather than promising universal compliance. See [Chapter 25](#) for the full legal posture.

21.1 Settlement as a First-Class Workload

The stack treats settlement itself as a **canonical workload**:

- Workload type Settlement(S, policy) includes chain(s) involved, transfer size, corridor routing, and policy hooks.
- Proofs at Layer 4 can attest that settlement followed policy and the final state matches expectations.

This means:

- Layer 5 isn't "just payments"; it is a **proven computation** over state transitions.
- Settlement events generate PIDL receipts that can be audited, archived, or collateralized.

21.2 Design Constraints for Privacy Rails

Privacy rails must satisfy a tight triangle:

1. **Non-custodial.** Users never relinquish both principal and unilateral control to intermediaries. Adaptor signatures, HTLC-like constructs, or scriptless scripts enforce this.
2. **Refund-safe.** If a swap fails, both sides can unilaterally reclaim funds after bounded time. Refund logic is testable and provable.
3. **Privacy-by-default with optional disclosure.** Default flows have strong anonymity sets; viewing keys allow specific flows to be revealed.
This is where the **Privacy Rails Kit (PRK)** lives.

Policy = Predicates, Not Surveillance

Policy compliance in Layer 5 is achieved through **predicate proofs**, not graph inspection.

What can be proved (ZK-compatible):

- Not on sanctions list → Membership proof: “sender $\in$ AML-cleared set S ”
- Within limits → Range proof: “amount $\leq$ \$10,000”
- Authorized counterparty → Set membership proof
- Tax reporting → Viewing key + receipt disclosure

What cannot be proved (and we don’t try):

- “Did not interact with blacklisted address X ” — requires global graph knowledge
- “Full transaction history disclosure” — defeats privacy model
- “Real-time surveillance feed” — centralizes and doxxes

21.3 The Privacy Rails Kit (PRK)

PRK is the Layer-5 sibling of PAL. It gives builders a way to express **pay-for-proof** and **pay-for-compute** intents and execute them over privacy-preserving, non-custodial corridors.

At a high level, PRK supports:

- **BTC $\leftrightarrow$ ZEC/XMR corridors.** Using adaptor signatures or HTLC-like patterns, ensuring atomicity. Prior art is directly deployed: the COMMIT-network Monero–Bitcoin atomic-swap construction and its open-source descendants (see Sources, “Privacy Rails”) are the operative reference for cross-chain atomicity over these pairs, and Chapter 21’s refund-safety invariants are stated so that such corridors can be certified against them rather than merely named.
- **Shielded pools and internal flows.** For in-asset privacy with proofs for policy compliance.
- **Conditional payment flows.** “Pay address A if proof P of workload W is posted within time T; else refund.”

PRK exposes intents like:

```
pay_for_proof(claim_id, max_price, corridor_policy)
execute_corridor(btc_input, zec_output, privacy_policy)
batch_payroll(payroll_blob, corridor_set, audit_policy)
```

21.4 Settlement Safety: Refund and Bridge Invariants

Non-custodial rhetoric is cheap; invariants are not.

Refund Safety Invariant: For any corridor C and transaction T :

$$\text{refundSafe}(T) = 1 \iff (T \text{ succeeds} \vee T \text{ refunds within timeout})$$

Where:

- Timeout is bounded (e.g., ≤ 24 hours for BTC $\leftrightarrow$ ZEC).
- Refund means unilateral recovery without counterparty cooperation.
- Success means both legs complete atomically.

Bridge Safety: Cross-chain bridges are historically the single biggest loss vector. Layer 5 imposes:

- **No “magic multisigs.”** Bridges must not rely solely on small, permissioned signer sets.
- **Light-client or proof-based validation where possible.**
- **Explicit trust classifications** reflected in profiles and telemetry.

Bridge Class	Trust Model	Collateral Eligibility	Example
Class A: Trustless	Cryptographic only	Full	Atomic swap
Class B: Threshold	m-of-n MPC	10% haircut	5-of-9 threshold
Class C: Federated	Known, bonded federation	25% haircut	Liquid-style
Class D: Custodial	Single custodian	Not eligible	CEX-wrapped

Table 21.1: Bridge trust classification.

Rule: Class D custodial bridges are excluded from SoV backing. Assets held via Class D bridges do not count toward collateral or reserve calculations.

21.5 VerifySettle: Settlement Telemetry

VerifySettle tracks:

- **Success rates** across corridors (p50/p95, by size bucket).
- **Refund outcomes** — how often refunds are exercised and their timing.

- **Anonymity-set health** — shielded pool sizes, churn, volume.
- **Concentration metrics** — top-N LPs, corridor operators, jurisdictional spread.

Metric	Definition	Target	Sev-1
swap_success(C)	% of swaps that complete	$\geq 95\%$	$< 90\%$ for 7 days
refund_safe(C)	% of failures that refund	100%	Any violation
ttf(C)	Time-to-finality (p95)	≤ 30 min	> 2 hours
anon_set(C)	Anonymity set size	≥ 1000	< 500 for 30 days
lp_concentration(C)	Top-3 LP share	$< 50\%$	$> 70\%$ for 14 days

Table 21.2: VerifySettle corridor health SLOs.

Corridor Health SLOs:

21.6 Layer-5 Stress Tests

Layer 5 earns the label “Value & Settlement” if it can navigate:

On-/off-ramp strangling. Can users still move value between BTC and privacy assets without centralized exchanges in country C ? Do non-custodial corridors maintain $\geq 95\%$ success with 100% refund safety?

Corridor LP exit. If a major liquidity provider disappears, do corridors collapse, or does routing adapt?

Regulatory attack on privacy. Can lawful-privacy corridors (viewing keys + receipts) keep enough demand to sustain anonymity sets when shielded assets are labeled “high risk”?

21.7 Minimum Viable Layer-5 Corridor

Scope Type	What Is Revealed	Who Holds Key	Use Case
Transaction-scoped	Single transaction details	Counterparty or auditor	Receipt for specific payment
Time-bounded	All transactions in window $[t_1, t_2]$	Auditor with consent	Tax audit
Amount-bounded	Transactions $>$ threshold	Compliance officer	AML monitoring
Full	All transactions	User only (default)	Personal records

Table 21.3: Viewing key disclosure scope model.

Viewing Key Disclosure Scope Model: **Key management:** Users generate and control all viewing keys. Keys are derived hierarchically; revoking a parent key invalidates children. No “master key” exists.

Requirement	Threshold
refund_safe	1.0
swap_success	$\geq 90\%$
ttf (p95)	≤ 2 hours
anon_set	≥ 500
Bridge class	A, B, or C only
Viewing key support	Transaction-scoped minimum

Table 21.4: Minimum viable corridor requirements.

Minimum Viable Corridor Spec: Corridors failing these requirements are flagged as “experimental” and excluded from default PRK routing.

Chapter 22

The Modular Stack

Parts III and IV described the stack layer-by-layer. This section flips the view: instead of layers, think in terms of **modules** builders actually touch and the **primitives** the system enforces underneath.

22.1 Primitive Catalog: The Twelve Pieces

Twelve primitives, and collectively they do one thing: they let two strangers trade verified work without either one trusting the other’s platform. That is the whole catalogue. When proofs from this kit clear, we mint verified FLOPs rather than heat; when the receipt schema binds, a claim becomes portable instead of platform-bound; when the corridors settle non-custodially, the value moves without a custodian appearing in the middle. No single primitive accomplishes any of that. The twelve are what the stack ships so that trust can be withdrawn, one counterparty at a time.

They fall into four clusters:

Compute / Consensus Primitives (AI-Money Substrate)

1. **MatMul-PoUW kit** — A work function that turns matrix multiplication into a useful lottery with verification asymmetry. Canonical sizes (e.g., MATMUL_4096), seeding rules, and adversarial tests ensure that:

- production cost scales $\sim O(n^3)$,
- verification cost scales $\sim O(n^2)$,
- and cheating is more expensive than being honest.

This is the Duplex-style substrate: when proofs from this kit clear, we mint **verified FLOPs** rather than heat.

2. **Verified-inference harness** — An Ambient-style proof-of-logits layer for AI workloads: deterministic transcripts, randomized audits, peer-prediction, and stake-and-slash for dishonesty. Where full ZKML is not yet practical, the harness gives us:
 - probabilistic guarantees that outputs came from model M under policy P ;
 - receipts per inference that VCO and PFS can route and price.

As ZKML matures, the same harness becomes the “front door” for full proofs. Until then,

it is what keeps an inference receipt from being a screenshot of a trust-me box: a transcript exists, an auditor can re-run it, and a prover who fabricates one loses stake.

3. **Canonical workload registry** — A dictionary of standardized SKUs for useful work: MATMUL_4096, INFER_LM_70B_256TOK, PROOF_2^20, PROVENANCE_VIDEO_V1, SETTLEMENT_BTC_XMR_V1, etc. This is the thesis’s unit of account in physical form: without it, “compute” is a negotiating position rather than a quantity, and no two quotes mean the same thing.

Inference Proof Tier Taxonomy Verified inference is the most economically significant—and currently most fragile—market under the “AI Money” analytical lens. The label does not make inference claims monetary; the guarantees are tiered explicitly so weak assurance is not marketed as strong:

Tier	Verification Type	Description	Error Bound
Tier A	Cryptographic correctness	Full ZK proof of inference (ZKML)	0 relative to the proved (quantized) computation; quantization
Tier B	Probabilistic soundness	Audited transcripts + randomized checks + peer prediction	Explicit (e.g., 99% confidence)
Tier C	Attestation-backed	TEE attestation + sampling audits	Implicit (trust TEE + sampling)

Table 22.1: Inference proof tier taxonomy.

Use Case	Minimum Tier	Rationale
Work Credit issuance (full weight)	Tier A	High-assurance service claims require cryptographic verification
Work Credit issuance (discounted)	Tier B	Acceptable with explicit error bounds and service haircut
Service market only	Tier C	Can sell inference; cannot mint high-assurance Work Credits
Collateral (pristine)	Tier A only	DeFi integrations require cryptographic certainty
Collateral (standard)	Tier A or B	With appropriate haircut
Collateral (ineligible)	Tier C	Not acceptable as collateral

Table 22.2: Tier eligibility rules.

Why this matters: Without this taxonomy, “verified inference” conflates cryptographic proof with probabilistic audit with vendor attestation. By making tiers explicit and tying them to Work Credit issuance weights, the thesis acknowledges engineering reality while preserving the instrument hierarchy: assurance grades service claims; monetary candidacy belongs only to the separate base asset.

Tier B requirements (reference): For Tier B verification to qualify:

- **Transcript reproducibility against a pinned kernel:** FP32 arithmetic is non-associative and BLAS kernels differ across hardware profiles, so bit-identical transcripts across heteroge-

neous provers are not achievable. The requirement is *pinned-kernel* reproducibility: each workload registry entry names a reference kernel (arithmetic order, FMA contraction off, no TF32), a transcript is reproducible by any prover running that kernel, and cross-kernel divergence is handled through the workload's relative error bound rather than exact-match. A prover may run a faster kernel only by proving conformance within ε_r of the reference result.

- **Audit rate:** $\geq 10\%$ of inferences are independently re-executed by random auditors.
- **Peer prediction:** Multiple independent provers; divergence beyond ε_r triggers investigation.
- **Stake-and-slash:** Provers stake collateral; fraudulent transcripts $\rightarrow$ 100% slash.
- **Error bound publication:** Explicit statistical guarantee (e.g., "99% confidence that output matches claimed model within ε_r ").

Migration path: As ZKML matures: (1) Tier A coverage expands to more model classes; (2) Tier B issuance weight increases toward $1.0\times$ as probabilistic guarantees tighten; (3) Tier C is deprecated from WC eligibility entirely. The goal is for inference verification to converge to Tier A over time, but the system functions with tiered guarantees in the interim.

Proof Primitives (Receipts as First-Class Objects)

4. **PAL SDK (Proofs-as-a-Library)** — Developer surface for Layer 4.
5. **Multi-ZK adapter & auto-selector** — Policy engine keeping proofs a commodity, not a vendor feature.
6. **PIDL (Proof Interface Definition Language)** — Minimal receipt schema binding claim, workload, proof, SLA, timestamps, and signatures.

Privacy & Settlement Primitives (Private-Money Substrate)

7. **Adaptor-signature atomic-swap kit** — Library for non-custodial BTC $\leftrightarrow$ ZEC/XMR settlement with refund-safety. The adversary here is the counterparty who takes your funds and disappears mid-trade, and the kit exists so that no one—including the corridor operator—can be that counterparty. Non-custodial rhetoric is cheap; invariants are not.
8. **Lawful-privacy corridor pack** — Hooks for viewing keys + receipts enabling regulated actors to prove compliance. A bank that cannot demonstrate compliance to its regulator cannot route anything, so these hooks are what keep lawful privacy lawful rather than merely private.
9. **Bridge-safety templates** — Pattern book for cross-chain settlement avoiding "magic multi-sigs." The graveyard this pattern book walks past is full of bridges whose security model was, on inspection, five signatures held by people who knew each other.

Market & Telemetry Primitives

10. **SLA escrow & slashing contracts** — Standard contracts for Bronze/Silver/Gold tiers with automatic refunds or slashes.
11. **VerifyPrice telemetry & methodology** — Measurement side of verification asymmetry.
12. **Neutral router & fairness tests** — Matching logic with house-share caps and entry-latency metrics.

PFS (Proof Factory Stack) bundles primitives 4–6, 10–12 to turn claims into routed proof jobs.

VCO (Verified Compute Orchestrator) bundles 1–3, 5, 10–12 to turn useful workloads into Work Credits.

22.2 The Four Reference Applications at a Glance

The four applications that follow are not a product roadmap. They exist to show that the stack closes: each one exercises a different span of layers, reports to a different set of public boards, and would trip a different set of red lines if the thesis is wrong. An architecture that cannot name which application would reveal which failure is not falsifiable in practice, however many metrics it publishes.

Application	Layers exercised	Reports to	Stresses	What its failure would show
Private treasury & payroll (Section 22.3)	L1 reach, L2 client integrity, L3 scoped claims, L4 policy proofs, L5 corridors	Settlement & Privacy; Agency Preservation	RL2, RL7, RL12	That lawful privacy is not operable at institutional scale—selective disclosure has become administrative disclosure
Media provenance (Section 22.4)	L0 attested capture, L2 distribution, L3 claims, L4 proof generation	Proof & Compute; Neutrality & Admission	RL1, RL3, RL10	That provenance cannot outlive the platform hosting it—receipts have become labels again
Verified inference (Section 22.5)	L0 capacity, L4 PoUW and proving, L5 contingent payment	Proof & Compute; Layer 0 Capacity; Economic Coverage	RL1, RL11, RL13	That verified compute earns no durable premium over ordinary cloud—the AI Money claim reduces to hosting
Proof/compute procurement (Section 22.6)	L4 workload registry, L5 settlement, L6 SLOs and telemetry	Economic Coverage; Value Capture; Wrapper Dominance	RL5, RL6, RL9	That demand was speculative rather than budgeted, and value leaks to operators and wrappers instead of the asset

Table 22.3: The four reference applications mapped to layers, public boards (Section 24.9), and the red lines (Section 28.6) each would trip first.

Between them the four applications stress eleven of the fifteen red lines. The four they do not reach—Red Line 4 (telemetry capture), Red Line 8 (governance capture), Red Line 14 (the capturable wedge closing), and Red Line 15 (the native collateral–capacity spiral)—are protocol-wide or market-wide failures rather than application-level ones, which is precisely why they cannot be caught by shipping an application and must be watched institutionally instead. Red Lines 14 and 15 are the clearest cases: an individual application can be technically successful while the capturable wedge closes or native collateral reflexivity degrades system-wide capacity.

22.3 Reference Application: Private Treasury & Payroll

It is the 26th of the month, and the payroll run is late. The company’s CFO is in Singapore, its controller is in Lisbon, and its 340 employees are spread across eleven jurisdictions. Under the incumbent way, the 26th is late because payroll is a daisy chain of NDAs: the bank wants the salary list, the auditors want the bank, the tax authority wants the auditors, and every one of those disclosures is a copy of the transaction graph that will never be recalled. The controller cannot prove the run is correct without handing over the graph, and cannot protect the graph without becoming the only person who can prove anything. Every quarter the company chooses which of those two failures to have.

This application removes the choice. Salary amounts stay confidential, the transaction graph stays private, and the auditors still get scoped evidence that every payment was authorized, complete, and properly reported—not the underlying ledger, but proofs about it, limited to what the mandate actually asks.

The flow is the stack, end to end:

1. **Create.** Finance defines a payroll batch with policy.
2. **Compute.** Application logic computes net pays and corridor selection.
3. **Prove.** PAL generates proofs that payments are within policy.
4. **Settle.** PRK executes payments over BTC↔ZEC/XMR corridors.
5. **Verify.** Auditors verify correctness via proofs and receipts.

22.4 Reference Application: Media Provenance

A clip surfaces on a Tuesday: a minister, a private conversation, forty seconds that would end a career if they were real. By Thursday the newsroom knows it is synthetic and cannot say so with anything it can hold up. The platform’s provenance label says “authentic”; the label is a string in the platform’s own database, and everyone in the room understands what that is worth. The newsroom’s correction cites “internal review.” The clip keeps circulating. Nobody lied; there was simply nothing to check.

This application is what there is to check. Capture hardware attests at the moment of recording, provenance proofs travel with the footage rather than living in a platform's database, and the question “did this come from camera C and reach us through the stated chain?” acquires an answer a reader can verify without trusting the newsroom, the platform, or anyone's internal review.

The flow:

1. **Create.** Cameras capture footage with Layer-0 attested hardware.
2. **Prove.** Proof factories generate Provenance(C, chain) proofs.
3. **Settle.** Advertisers pay creators over privacy rails, conditioned on valid proofs.
4. **Verify.** End-users verify that content passed through the stated chain.

22.5 Reference Application: Verified Inference

A risk officer at an insurer signs off on a model whose outputs will move payouts on thousands of policies. The vendor's API returns a number and a confidence score, and behind the endpoint the officer can see nothing: not whether the model is the one advertised, not whether the input was modified in transit, not whether the “confidence” is a measurement or a marketing decision. The incumbent answer is a due-diligence visit and an attestation letter. The officer files the letter and assumes residual risk, because there is nothing else to do with it.

This application replaces the letter with a receipt. The enterprise submits an inference request under stated constraints; the proof factory runs the model and produces Inference(M , X ; policy) with a proof that the output came from model M on input X under that policy. Payment is contingent on the proof being valid, so the vendor is paid for verified work, not for uptime of a black box.

The flow:

1. **Create.** Client submits inference request with constraints.
2. **Compute.** Proof factory runs the inference.
3. **Prove.** PAL compiles into Inference(M , X ; policy) with proof.
4. **Settle.** Payment via PRK contingent on valid proof.
5. **Verify.** Client verifies proof and accepts output.

AI Money is an analytical lens or legacy label for these inference-service markets. The instruments are typed Work Credits and derivatives; they are not themselves money.

22.6 Reference Application: Proof/Compute Procurement

A DAO's treasury committee is staring at a proof-capacity crunch it can see coming and cannot buy ahead. Under the incumbent way, the committee's options are a cloud commitment letter—a price the vendor sets, hardware the committee cannot inspect, and a lock-in it will

still be paying for after the crunch passes—or forward contracts on capacity nobody can verify exists. Neither instrument lets the DAO prepay for capacity it can later prove it received. The gap is not a pricing gap; it is a verification gap, and it is why procurement at distance has historically meant procurement on faith.

This application closes the gap. The DAO declares demand curves for canonical workloads over time, buys typed Work Credits as they are minted against verified work, and holds them as prepaid service inventory. Anyone can audit the backing: the credits are minted against registry workloads whose proofs and VerifyPrice profiles are public, so the committee's claim to future capacity is checkable by any token holder, not asserted by the committee.

The flow:

1. **Create.** DAO defines demand curves for workloads over time.
2. **Compute.** Procurement module computes optimal schedules.
3. **Prove.** As Work Credits are minted, DAO purchases them.
4. **Settle.** Payments over privacy rails; typed Work Credits held as prepaid service inventory, not monetary reserves.
5. **Verify.** Anyone audits backing via VERIFYPRICE, profiles, and proof metadata.

The DAO's treasury holds BTC/ETH reserve assets beside typed proof and compute capacity instruments. "ZK Money" and "AI Money" are analytical or legacy lenses for those service markets, not claims that the underlying Work Credits are money. Any separate base asset remains only a conditional monetary candidate.

22.7 Part IV Summary

Part IV is where the triad learns to **speak economics**:

- **Layer 4** turns computation and provenance into standardized, verifiable work units.
- **Layer 5** turns those work units into private, non-custodial value flows.

With these in place, the legacy **Private Money**, **AI Money**, and **Proof/Attestation Money** labels become useful analytical lenses on settlement, compute, and proof-service demand. Their Work Credits and capacity instruments remain typed service claims, not monetary instruments. Only the separate base asset may become a monetary candidate if all conditions pass. In Part V, we climb to **Layer 6 – Governance & Telemetry**, where neutrality and repression-resilience are kept falsifiable, not just promised.

Part V

Governance, Telemetry & Neutrality

Parts I–IV argued that the old soft guarantees are failing; that Privacy, Proofs, and Compute are distinct services that may support a conditional base-asset monetary candidate; that a seven-layer cypherpunk stack can supply these services; and that Layers 4–5 turn work into receipts and receipts into private, non-custodial settlement.

Part V looks at the same system from the **operations and governance** angle. If the base asset is to remain a store-of-value candidate under repression, the stack’s service delivery and holder-side agency must be measured rather than promised. Work Credits remain typed service claims; their usage cannot substitute for native monetary evidence.

Chapter 23

Layer 6: Governance & Telemetry

Layer 6 is the **control plane** and the **nervous system** of the stack.

Everything below it—hardware, comms, distribution, identity, proofs, settlement—can drift, centralize, or quietly break without anyone noticing until it’s too late. Layer 6 refuses to let that drift remain invisible. It insists that:

- **Neutrality** (no favored flows, no hidden house edge),
- **Repression-resilience** (still works under capital controls and censorship), and
- **Verification economics** (VERIFYPRICE, VerifyReach, VERIFYSETTLE)
are **service-level objectives (SLOs)**, not marketing copy.

23.1 “No Dashboards, No Trust”: Public SLOs as Constitution

Most protocols ship documents called “constitutions.” In practice, the real constitution is whatever you cannot silently violate without getting caught.

For this stack, the constitution is:

- a **set of SLOs** on triad capacity and neutrality, plus
- a **set of dashboards and receipts** that make violations obvious.

Examples:

- “p95 VERIFYPRICE for canonical workloads stays below T seconds on commodity hardware.”
- “Top-N prover/LP/router share remains below $X\%$ by hashpower/FLOPs/liquidity.”
- “BTC↔ZEC/XMR corridors achieve $\geq 95\%$ success with 100% refund safety.”
- “At least K distinct jurisdictions and hardware profiles are actively verifying.”

These are not just operational goals; they are the **monetary invariants**. Layer 6’s first job is to **publish and maintain these SLOs as public contracts**.

23.2 Control Surfaces: Parameters, Upgrades, Emergency Powers

Every protocol has knobs. The question that decides whether it is a constitution or a clique is not *which* knobs exist but *who can turn them, with what evidence, and how visible the turning is*. The second job of Layer 6 is to inventory the control surfaces — to name the knobs, then bind each one to a proposer, a ratifier, a data requirement, a delay, and a rollback path. A knob that is not on this list is not governed; it is somebody’s private steering wheel.

The knobs come in three families. **Economic parameters** are the ones that touch money directly: block reward curves and Work Credit issuance schedules, fee policies (floor, burn, and the split between security and dividends), and PoUW weighting across workloads. **Technical parameters** shape what the network will accept: circuit versions and proof-system parameters, additions and deprecations to the accepted hardware profiles, and corridor policies such as timelocks, refund windows, and anonymity thresholds. **Operational parameters** govern the watching itself: telemetry granularity and reporting requirements, incident severity levels and response playbooks, and key ceremonies with their quorum sizes. The families differ in blast radius, which is exactly why they are enumerated separately rather than dumped into one list — issuance is not incident management, and pretending they are the same class of decision is how emergency powers leak into monetary policy.

For each surface, Layer 6 specifies: (1) who can propose a change, (2) who can ratify it, (3) what data must be presented, (4) what delays and rollback mechanisms exist, and (5) which changes are “emergency powers.”

23.3 Bell-Labs-Style R&D vs. On-Chain Governance vs. Off-Chain Norms

There are three governance “forces” that need to be reconciled:

1. **Bell-Labs-style R&D.** A research org with autonomy to explore new circuits, proof systems, hardware profiles, and corridor designs.
2. **On-chain governance.** Token- or Work-Credit-weighted voting, parameter changes baked into protocol logic.
3. **Off-chain norms and institutions.** Foundations, labs, dev collectives, and community norms.

Layer 6’s thesis is not “pick one.” It is:

Use Bell-Labs-style R&D to discover, on-chain governance to ratify and constrain, and off-chain norms to fill the gaps—but keep all three under telemetry.

Concretely:

The R&D lab:

- Maintains public roadmaps and risk registers.
- Publishes upgrade proposals with quantified impacts on VERIFYPRICE, VerifyReach, VERIFYSETTLE, decentralization, and hardware profiles.
- Runs testnets and shadow deployments.

On-chain governance:

- Controls **scarce resources**: issuance, reward splits, canonical workloads, acceptance/deprecation of proof systems and hardware profiles.
- Is bounded by **constitutional SLOs**: certain changes are simply not allowed if they push metrics beyond thresholds unless higher-order safety processes trigger.

Off-chain institutions:

- Operate under **public charters** that explicitly state their mandate and constraints.
- Are expected to publish minutes, risk assessments, and incident reports.

This three-body system is inherently unstable; Layer 6 keeps it from drifting into pure plutocracy or pure priesthood by insisting that **all three bodies are visible in telemetry**:

- Changes in code and parameters appear on-chain.
- Lab work appears in open repos and benchmarks.
- Institutional decisions appear in charters, public calls, and reports.

No dark corners; no “just trust us, we’re the stewards.”

Governance Capture Risk

Token-weighted governance creates capture risk: large holders can centralize control over issuance, fee routing, and telemetry rules. The stack mitigates this through hard constitutional constraints that governance cannot override (see [Section 23.5](#)), timelocks on parameter changes, independent measurement infrastructure, and non-governable red lines ([Section 28.6](#)). If governance can alter issuance, fee routing, or telemetry rules without hard constraints, this is itself a red line.

23.4 Governance as SLOs: The Five Invariants

Governance and operations exist to keep one invariant true under stress:

Pay the machine only for work anyone can verify cheaply—and give humans lawful privacy by default.

Everything below—roles, metrics, runbooks, and legal posture—turns that sentence into procedures, with receipts. Where policy pressure rises, we default to measurement and non-custodial design rather than new gatekeepers. “No dashboards, no trust” is not a slogan; it is the rule for deciding when the system is still safe to treat as money.

Governance and operations are about what happens when the world pushes back. They are the difference between a beautiful mechanism that works in a friendly lab and an actual monetary substrate that survives YCC (yield-curve control), capital controls, app-store bans, and “secure enclave” mandates.

We can summarize the protocol’s “constitution” as **five invariants**:

1. **Verification asymmetry.** For each canonical workload W , verification stays much cheaper than production. Formally, the verification overhead $r(W) = v(W)/p(W)$ remains at or below a published bound (e.g., 0.3, with a stretch goal of 0.1), and this is reported via $\text{VERIFYPRICE}(W)$ on the Proof & Compute Board. If $r(W)$ drifts up or p95 verify times blow past targets, the economic hinge of the triad fails; this is a Sev-1 condition.
2. **Open admission.** Anyone who brings correct work clears. Admission is not gated by identity, licensing, or proprietary hardware. Routers are neutral and open-source; house share and time-to-first-fill for new provers/LPs are visible on the Neutrality & Admission Board. If honest new entrants cannot join and get paid within a bounded window, decentralization is already drifting.
3. **Useful-work security budget.** Block rewards and fees underwrite **useful work**—MatMul-PoUW, verified inference, provenance and settlement proofs—under explicit SLAs. The security budget is tied to canonical workloads with published VERIFYPRICE rather than to a puzzle with no external buyer. That is a deliberate trade and not a free improvement: a work function with a buyer has a demand curve somebody can move, which is the objectivity a buyerless hash puzzle gets for nothing (Section 31.1.1).

Inference Verification Tiers and WC Eligibility: Under the “AI Money” analytical lens, inference Work Credits remain service claims. To prevent weak spot checks from being marketed as high-assurance verification, workloads are tiered by verification strength:

Tier	Verification Type	WC Eligibility
Tier A	Cryptographic (full ZK)	Full eligibility (1.0×)
Tier B	Probabilistic (bounded error, audited)	Discounted (0.6×)
Tier C	Attestation-only (TEE + sampling)	Service market only; no WC issuance

Table 23.1: Inference verification tiers and WC eligibility.

Rule: Only Tier A and Tier B verified inference is eligible for Work Credit issuance. Tier B issuance is discounted by a published *issuance weight* (the $0.6\times$ of the tier table in [Section 20.10](#)) reflecting error bounds and audit rate — a weight on what is minted, not a haircut on what is pledged. Tier C can exist as a service market but is not collateral-grade—it represents trust in TEE vendors, which is exactly what the thesis aims to minimize.

4. **Lawful privacy by design.** Settlement is neutral and non-custodial by default—adaptor-sig atomic swaps, shielded pools, privacy rails—but comes with **viewing keys and auditable PIDL receipts** so law-abiding users can evidence obligations without re-introducing choke-points. The Settlement & Privacy Board reports swap success, refund safety, and anonymity-set health; flows that never touch custodians should still leave enough receipts that auditors can do their jobs.
5. **Verifiable machines at Layer 0.** Canonical proving and attestation paths run on open designs with sampled supply chains; hardware attestation becomes one input to proofs, not a vendor priesthood. Hardware profiles, lot-sampling coverage, and profile incidents show up on the Hardware/Layer-0 panes of the boards.

These invariants are **hard to change and easy to measure**. Everything else—fee curves, circuit versions, work-function parameters, corridor lists—is configuration.

To keep configuration from degenerating into a governance circus, we tie changes to **observable triggers** instead of vibes:

- If **VERIFYPRICE p95** for a workload W drifts above its target band for a sustained window, that workload’s circuit and parameters are queued for revision. Blocking the change requires an explicit override that cites alternative SLOs and appears in the governance log.
- If **entry latency** for new provers/miners/LPs rises beyond a threshold, or the **top-N share** and house share breach caps, neutral routers automatically ratchet down house share and prioritize small bidders until the metric recovers.
- If a **swap corridor’s refund-safety** ever falls below 100% in VERIFYSETTLE, that corridor is removed from admissible routes by default. Re-listing requires a synthetic regression suite to pass (including stress tests) plus a public post-mortem.

In this frame, “governance” does not micromanage every adjustment; it **writes the SLOs and the triggers**. Operations enforces them and publishes the receipts. Most changes are **data-driven roll-forward**: parameters evolve in response to Verify* drift, with clear before/after numbers.

To keep rule-making separate from rent-collection, we distinguish two broad roles:

- A **spec & telemetry steward** that publishes ABIs, PIDL schemas, canonical workloads, hardware profiles, and the VERIFYPRICE/Reach/Settle dashboards. It does not run routers or take custody. Its mandate is to keep the bill of materials and metrics honest and up to date.
- **Market participants**—miners, provers, routers, corridor LPs—who compete on price and reliability under those public SLOs, with SLA escrow and slashing keyed to PIDL receipts.

Their incentives are economic; their constraints are the invariants and metrics.

This “two-chamber” structure keeps the job of defining rules and SLOs distinct from the job of selling capacity under those rules. It avoids the familiar pattern in which the entity that can edit parameters also happens to own the biggest fleet of nodes.

Governance as SLOs, not personality, is what lets the system answer policy pressure with dashboards and runbooks instead of press releases. When asked “how do you survive X?”, the right answer is not “trust the foundation,” but “look at this board, this trigger, and this incident playbook.”

23.5 Constitutional Enforcement

“SLO-bounded governance” is only credible if the enforcement mechanism is specified.

23.5.1 Machine-Enforced Constraints (Hard)

Certain invariants are enforced **on-chain or in protocol code**:

Constraint	Enforcement
Issuance cap per epoch	Smart contract rejects minting transactions exceeding cap
Corridor refund timeout	Atomic swap contracts enforce timeout
House-share cap	Neutral router contracts reject bids exceeding share
VERIFYPRICE bound violation	Workload suspended from WC eligibility

Table 23.2: Machine-enforced constitutional constraints.

23.5.2 Override Path (Slow, Expensive)

Some constraints need flexibility. The override path is **deliberately expensive**:

- **Supermajority:** $\geq 67\%$ of governance weight.
- **Long timelock:** ≥ 30 days (90 days for issuance-related).
- **Risk memo:** Mandatory written analysis, hash anchored on-chain.
- **Sunset clause:** Override auto-expires after 12 months.

23.5.3 Emergency Path (Narrow)

True emergencies (proof system break, active exploit) require faster action. The emergency path is **narrow and auditable**:

Requirements: ≥ 3 of 5 security signers; auto-expire in 7 days; postmortem within 14 days.

Allowed	Not Allowed
Deprecate proof system	Increase issuance
Delist corridor	Confiscate user funds
Quarantine hardware profile	Bypass refund safety
Freeze workload class	Mint unbacked credits

Table 23.3: Emergency path scope limitations.

Reconciling the emergency path with Red Line 8. Section 28.6 retires the thesis if governance can alter issuance, fee routing, or telemetry rules without hard constraints. Read carefully, this path appears to grant exactly that: five signers on a seven-day fuse. The reconciliation is structural and worth stating as a constitutional rule rather than leaving to inference:

- **Emergency powers are subtractive only.** Every allowed action removes or restricts a protocol claim (deprecate, delist, quarantine, freeze). No allowed action redirects value: no fee re-routing, no issuance change, no collateral-rule change, no telemetry amendment. A body that can only shrink the protocol cannot steal from it; the worst case is a smaller protocol, which is a survivable failure and the correct bias for a seven-day fuse.
- **Subtractive is not free.** Quarantining a hardware profile concentrates verification (Red Line 3’s condition); freezing a workload class cuts its fee flows. Emergency use therefore carries its own falsification weight: emergency actions must be logged against the red-line conditions they aggravate, and *repeated or routine* use of subtractive powers — the “emergency governance becoming routine” failure named under Red Line 12 — trips Red Line 8’s spirit even though each individual action was constitutional. The test is not whether any single action was permitted but whether the emergency path has become a de facto governance channel.
- **Signers are named, accountable, and rotating.** The five seats are filled by published charter, occupied by independent security researchers and operators (not the foundation, not the largest holders), subject to rotation, and removable by the standard override path. Anonymity of the signer set is itself a Red Line 4 telemetry failure.

23.6 Monetary Constitution

23.6.1 Issuance Envelope

Work Credit issuance is governed by:

- **Base schedule:** Halving every N blocks (e.g., 4-year halvings).
- **Capacity modulator:** $\pm 10\%$ based on verified capacity growth.

Capacity modulator formula:

$$\text{Issuance}_{\text{epoch}} = \text{BaseSchedule}_{\text{epoch}} \times (1 + 0.1 \times \tanh(\text{CapacityGrowthRate} - \text{TargetGrowthRate}))$$

23.6.2 Fee Routing

Destination	Share	Mechanism
Capacity providers	70%	Paid via SLA escrow
Protocol burn	20%	Permanently removed
Assurance budget	10%	Funds sampling, audits, security

Table 23.4: Fee routing split.

Why burn? Burn creates deflationary pressure tied to usage. High demand → high fees → high burn → supply reduction → value appreciation. This closes the loop between utility and asset value.

This routing is what makes the [Value Capture Lemma](#) from [Chapter 7](#) more than a slogan: it is the fee plumbing that delivers the lemma’s cash-flow claim. Its 20% burn is the same conservative variant used in the Layer 4 worked example ([Chapter 20](#)) rather than the 30–50% reference band of [Chapter 7](#), because here the remaining 10% funds assurance. The burn component is what converts usage into scarcity; the assurance budget funds the Layer 0 sampling that keeps hardware claims credible. Changes to the fee split are governance parameters, but reducing the burn below 10% or eliminating it entirely would undermine the value-capture loop and should be treated as a constitutional constraint.

23.6.3 Retirement Rule

In addition to fee burn, base-asset units are retired in specific circumstances:

Trigger	Retirement
Work Credit redemption	Any base-asset fee paid to fulfill the capacity claim is retired under the fee rule; the Work Credit itself is retired as a fulfilled service claim
Slashing	100% of slashed collateral is burned (not redistributed)
Failed workloads	Fees for failed proofs (prover fault) are burned, not paid

Table 23.5: Base-asset and service-claim retirement triggers.

23.6.4 Risk Haircuts

Work Credits minted under weaker assurance levels receive discounted issuance:

Factor	Haircut	Example
Hardware profile (L0 grade)	L0-A: 0.9×; L0-B: 1.0×; L0-C: 1.1×; L0-D: 1.2×	Lower-assurance hardware earns fewer credits
Verification tier	Tier A: 1.0×; Tier B: 0.6×; Tier C: 0.3×	Probabilistic verification earns less than cryptographic
Prover concentration	>20% share: 0.9×; >30%: 0.8×	Concentrated provers earn progressively less
FCR confidence grade	Missing/unverifiable FCR: issuance cap; low-confidence FCR: 0.85×	Facilities that cannot substantiate physical capacity claims earn fewer, capped credits
Sovereign optionality ($\mathcal{O}_s$)	High band: 1.0×; moderate: 0.9×; low: 0.75× and issuance cap	Facilities whose physical capacity is fully substantiated

These haircuts are protocol parameters (not discretionary); they appear in dashboards and apply automatically.

Two distinct physical factors. The last two rows look similar and do different work. The distinction matters enough to state explicitly, because collapsing them leaves a gap that a well-run fragile operator can walk through.

- **FCR confidence grade is epistemic.** It asks whether we can see the physical substrate at all. It penalizes opacity.
- **Sovereign optionality is strategic.** It asks whether what we can see is survivable. It penalizes fragility (Section 15.8).

Without the second row, a facility that reports complete, signed, independently verified FCR data documenting a single interconnect, one fuel source, no backup duration, and a jurisdiction with an active rationing regime would receive *no haircut whatsoever*. It would in fact score well, because its disclosures are impeccable. Transparency about fragility is not the same as resilience, and an issuance schedule that rewards only the former is measuring candour rather than durability.

This is the operational form of the distinction drawn in Section 15.10.1: verifiability and availability are separate properties, and Work Credit issuance should price both.

Interaction with the affordability hinge. These haircuts must not become a back door for degrading verification cost. Per Section 15.10.3, a resilience premium is acceptable only while benign-state Physical VerifyPrice stays inside its constitutional SLO band. If pursuing a higher $\mathcal{O}_s$ band pushes a network's verification cost outside that band, the SLO wins and the optionality target is deferred. Haircuts adjust issuance; they do not license breaching a constitutional target.

23.6.5 Coverage Target

The **fee-coverage ratio** measures what share of the security budget comes from fees vs. issuance:

$$\text{FeeCoverage} = \frac{\text{RetainedFees}}{\text{TotalSecurityBudget}}$$

where RetainedFees is fee revenue net of the burned share — the portion that actually funds validators, provers, and operations. Burns do not fund security: a burn is supply arithmetic, not security revenue (Appendix H.14), so adding it to the numerator would double-count the burned slice of fees (it is already inside gross fee revenue) or, if read as destruction of issuance, would count destruction as funding. The burned quantity is published as its own series and read as reduced net issuance. Gross fee revenue is likewise published, so the burn share is visible; the convention is pinned here and applies wherever a fee-plus-burn quantity appears


```

|      v      |
| NET EFFECT: Demand -> Fees -> Burns + Staking -> Accrual |
| Monetary premium still requires holder and agency conditions |
+-----+

```

Why demand doesn't simply expand supply proportionally:

1. **Issuance cap is hard:** Protocol enforces epoch-level cap regardless of demand.
2. **Burn is automatic:** Higher demand → higher fees → higher burn → lower net supply.
3. **Staking locks supply:** More provers/routers → more collateral locked → less circulating supply.
4. **Haircuts constrain issuance:** Weak verification or concentrated provers earn less, limiting supply growth.

This is the “bond indenture” that makes the SoV claim auditable, not asserted.

Chapter 24

Extended Telemetry

A constitution nobody can check is a press release. Layer 6 claims that measurement is the constitution — that SLOs, not assurances, decide whether the asset is still money — and that claim is only as strong as the measurement regime beneath it. Dashboards can be gamed, and the gaming is usually boring: a favorable region here, a fast machine there, a bad month quietly absent. So the observatory must be built the way it audits others. Anyone can re-run the audit, on hardware they own, and catch the observatory lying. That is the standard this chapter engineers for, and everything in it — reference machines, cost vectors, sampling seeds, anchored archives — exists to make that re-run possible.

24.1 Reference Verifier Classes

“Anyone can verify” is a promise about hardware, so it has to be priced on hardware a person can actually own. All Verify* metrics are therefore measured on **standardized reference hardware**:

Class	Hardware Spec	Use Case
Laptop-Class	4-core CPU, 16GB RAM, SSD, no GPU	Default for SLOs
Mobile-Class	ARM SoC, 8GB RAM	Edge/mobile
Server-Class	32-core CPU, 128GB RAM	High-throughput

Table 24.1: Reference verifier classes.

Rule: All constitutional SLOs (e.g., “VERIFYPRICE $t_{95} \leq 5s$ ” — the p95 *time* field of the canonical tuple, [Appendix A](#)) are defined against **Laptop-Class** unless a workload’s registry entry binds it to another class, in which case the entry names the class. This ensures the “anyone can verify” promise is falsifiable. Where a tier table quotes a different number for a named variant of the same workload (e.g. INFER_LM_7B_512TOK at 10s), the registry entry, not this blanket rule, governs the variant; the blanket rule governs the class.

Requirement	Specification
Independent implementations	≥ 2 independent verifier codebases per canonical workload tier
Open source	All reference verifiers must be open-source with reproducible builds
Versioned and hashed	Binary hashes are part of the canonical workload registry
Deterministic	Same proof + same verifier version $\rightarrow$ same result, always

Table 24.2: Reference verifier implementation requirements.

24.2 Reference Verifier Implementations

Why two implementations rather than one trusted one? Because a single verifier codebase is a single point of failure for the truth itself. One bug can inflate or deflate every measurement the constitution depends on, and nobody outside would ever know. Two independent implementations, built by different people from the same specification, turn that silent failure into a visible disagreement — and a disagreement, unlike a bug, shows up on a dashboard.

24.3 Cost Vector Definition

A single number is easy to quote and easier to flatter. VERIFYPRICE is deliberately not one. It is a **cost vector** for each workload W and verifier class V :

$$\text{VerifyPrice}(W, V) = (t_{50}, t_{95}, m, b, e, f)$$

Where:

- t_{50}, t_{95} : Median and 95th-percentile verification **time** (seconds)
- m : Peak **memory** usage (MB)
- b : **Bandwidth** consumed (bytes)
- e : **Energy** estimate (Joules)
- f : **Failure rate**

24.4 Sampling and Anti-Cherry-Pick

The most common attack on telemetry is not fabrication; it is curation. Nobody has to invent numbers when selective reporting will do: show receipts from good regions, fast hardware, favorable conditions, and the dashboard flatters itself without a single lie in it. The sampling regime exists to make curation detectable:

Rule: Any dashboard point published without a verifiable sampling seed and stratification breakdown is not part of the official Verify* record.

Mechanism	Specification
Public randomness beacon	Receipt selection seeded by on-chain randomness (e.g., block hash, VRF output)
Stratified sampling	Samples drawn proportionally from: workload class, hardware profile, geographic region, time-of-day
Adversarial corpus	10% of test proofs are malformed or worst-case (max witness size, pathological inputs)
Continuous measurement	Rolling 24-hour windows; no periodic snapshots that can be gamed
Multi-operator	≥ 3 independent measurement operators; divergence $> 5\%$ triggers investigation

Table 24.3: Sampling mechanisms for Verify* telemetry.

24.5 Data Availability and Anti-Memory-Hole

Dashboards are views. The source of truth is the **receipt corpus**. Receipts must be:

Requirement	Specification
Content-addressed	Receipt sets identified by Merkle root
Anchored	Merkle roots periodically committed to a neutral ledger (e.g., Bitcoin, Ethereum)
Archived	Raw receipt data retained for ≥ 2 years by ≥ 3 independent archivists
Queryable	Any observer can request receipts corresponding to a dashboard point

Table 24.4: Data availability requirements.

A dashboard can be edited; a Merkle root committed to a neutral ledger cannot be uncommitted. That asymmetry is the whole defense. If a government or a CDN can quietly delete the bad months, the telemetry regime has been captured without a single headline. Content-addressing plus neutral anchoring does not prevent deletion — it makes deletion *visible*, and visible deletion is itself a datum.

24.6 Reproducibility Contract

A valid dashboard datapoint is defined as:

$$\text{Datapoint} = (d_{\text{root}}, v_{\text{harness}}, v_{\text{verifier}}, r)$$

Validity check: Anyone can:

1. Fetch receipts by d_{root} (dataset root)
2. Run v_{harness} (harness version) with v_{verifier} (verifier version)
3. Verify r (result) matches

If the results don't match, the datapoint is invalid and the discrepancy is a Sev-1 incident.

Published artifacts:

- Harness code (open source, tagged)
- Verifier binaries (reproducible builds, hashed)
- Dataset roots (timestamped, anchored)
- Aggregation scripts (deterministic)

This transforms “no dashboards, no trust” from philosophy into a **reproducible measurement contract**.

24.7 How to Measure: Receipts, Not Vibes

VERIFYPRICE was introduced at Layer 4; VerifyReach at Layer 1; VERIFYSETTLE at Layer 5. Layer 6 treats them as a **coherent observability regime** rather than three unrelated graphs.

The observability story only works if it is itself verifiable. That is why every claim, proof, settlement, and SLA outcome is turned into a **receipt** with enough structure that anyone can replay the metrics on a reference verifier.

In practice, this means:

- Every proof receipt carries not only the proof artifact but also start/end timestamps, resource usage, and a hardware profile tag.
- Settlement receipts encode corridor, timing, refund status, and failure codes.
- Routers emit anonymized samples of order books, matched and unmatched bids, and house vs. third-party flags.
- Hardware vendors publish lot-attestation artifacts that receipts can reference as part of their profile.

A public “VerifyPrice Observatory” then ingests these receipts and emits the aggregate vectors we refer to in the text. Anyone can pull the raw receipts corresponding to a dashboard point, feed them into the same open-source verifier suite, and see whether they obtain the same numbers. If they do not, the mismatch is itself an incident.

This architecture enforces a simple discipline: **if it matters, there is a receipt for it**. If a claim about performance, neutrality, or safety cannot be tied back to receipts that anyone can re-check, it is out of scope for the monetary thesis.

They are the three primary dials for answering:

- *Is verification still cheap?* (VERIFYPRICE)
- *Can users still reach the network under pressure?* (VerifyReach)
- *Can users still move value privately and safely?* (VERIFYSETTLE)

We can think of this observability regime as watching the stack across five planes:

Proof & compute plane Are receipts reliable, and does verification asymmetry actually hold under load?

Settlement & privacy plane Do non-custodial payouts actually clear, with privacy and refund-safety matching the promise?

Neutrality & admission plane Can new provers, miners, routers, and mirrors join on equal footing, or has the system ossified into a club?

Economic coverage plane Are real fee flows, not issuance, paying a growing share of the security and operations budget?

Layer-0 hardware plane Do “verifiable machines” remain falsifiable in practice, or did we slide quietly back into “trust the vendor”?

Each plane maps directly to the SoV requirements from Part II: credible scarcity, cheap public verification, censorship-resistance, neutrality, native demand, lawful privacy, and duration-neutrality. If any plane drifts too far from its targets, the monetary thesis weakens, no matter how elegant VERIFYPRICE looks in isolation.

24.8 VerifyFlow: Instrumenting the Market Realization Plane

VERIFYPRICE, VerifyReach, and VERIFYSETTLE answer whether the stack is real. None of them answers a question that will dominate every public conversation about the asset: *why is the price doing that?* Corollary 11.1 says we cannot infer monetary adoption from price, and Section 11.14 says the answer lives in external market structure. VerifyFlow is the fourth named verification family, and it measures the external financial representation and price-transmission state of the native monetary object.

The governing principle extends the existing one:

No protocol dashboards, no monetary trust. No flow ledger, no price inference.

24.8.1 Wrapper-Level Vector

For each external financial wrapper i , publish

$$\text{VerifyFlow}_i = (L_i, A_i, \Delta A_i, \varepsilon_i, \kappa_i, C_i, H_i, D_i)$$

where L_i is the leverage multiple and reset rule; A_i is assets or delta-adjusted exposure; ΔA_i is creations and redemptions, published *raw and flagged*: ΔA_i is the input to the decomposition of Appendix H.1, not an aggregate in its own right, and no downstream sum may consume it unless the wrapper’s series has been decomposed into its ε -, α -, and residual components — otherwise aggregating across wrappers silently rebuilds the double-count the asset-level vector exists to prevent. r_i in the Q^{RC} term below is the wrapper’s underlying return over the daily-reset horizon; ε_i is holder flow elasticity; κ_i is net mechanical gain; C_i is constituent concentration; H_i is hedge implementation (physical, swaps, futures, options); and D_i is dealer or counterparty concentration. Appendix H derives ε and κ and states their measurement

contract.

24.8.2 Asset-Level Vector

For the native asset X , publish

$$\text{VerifyFlow}(X) = (Q^{RC}, Q^{RD}, \text{MPR}^{\text{gross/net}}, \text{CCR}, \text{SER}, \text{WNG}, \text{RCR}, \text{NUS})$$

The MPR slot is a published pair (gross and net), so the vector names the pair rather than one form; every other slot is a single series.

Return-coupled demand Q^{RC} . Exposure that moves as a function of the return itself, $Q^{RC} = \sum_i \kappa_i A_i r_i$. This term shapes volatility, momentum, persistence, and reversal severity.

Return-decoupled demand Q^{RD} . Allocation flow largely independent of the day's return, $Q^{RD} = \sum_k w_{Xk} \alpha_k A_k$, where α_k is the return-independent intercept of net creation into wrapper k and w_{Xk} its delta-equivalent exposure weight to X . Q^{RD} consumes α only. Per the decomposition rule of [Appendix H](#), creations split into an ε -driven component (return-coupled, already inside κ), the intercept α_k (structural bid — distribution growth, mandate adoption, scheduled contributions), and a residual u_k (unexplained, published as its own series and consumed by no aggregate). Feeding raw F_k into this sum would double-count the ε -component, which κ already carries through Q^{RC} , and would inflate MPR and every downstream reading built on it. This term shapes destination, concentration, and level — which prices get accepted without valuation-sensitive selling.

Mechanical Pressure Ratio (MPR). Gross measures *pressure*; net measures *direction*. That distinction is the whole metric, and it is worth the two definitions it costs. Q^{RC} and Q^{RD} partition product-level demand by return-dependence: the ε -component of creations belongs to Q^{RC} (inside κ) and only α_k belongs to Q^{RD} . Total product-level *exposure demand* is then $Q^{\text{exp}} = Q^{RC} + Q^{RD}$ with no term counted twice. Dealers translate that exposure into market orders, so realized market demand is its intermediated image, $Q^{\text{mkt}} = \mathcal{H}(Q^{\text{exp}}, \text{dealer inventory, internalization, swap})$. The gross form, $\text{MPR}^{\text{gross}} = \sum |Q_j^{\text{mkt}}| / V_X$, counts every dollar of liquidity consumed on both sides: it asks how hard the machinery is leaning on the market, whichever way it leans. The net form, $\text{MPR}^{\text{net}} = |Q^{\text{mkt}}| / V_X$, nets the offsetting flows first and asks only where the weight ends up. Here V_X is the asset's executable liquidity — average daily traded USD volume on the reference venue set over the same window the numerator aggregates, published with the ratio — and the segments j are the pre-declared reporting partition (venue $\times$ wrapper class), also published with the ratio.

The two forms must never be averaged, ranked, or otherwise combined, because each guards against a different deception. Netting offsetting flows before the modulus erases the liquidity both sides consumed, which is why the gross form leads and the net form contextualizes. And customer exposure and dealer hedging are the same demand observed at two stages, not two

independent demands; summing them would double-count. Report Q^{exp} and Q^{mkt} separately and never add them — a pressure gauge and a compass are two instruments, and gluing them into one dial breaks both.

Custodial Control Ratio (CCR). Base asset controlled by custodians and wrappers, divided by circulating supply.

Synthetic Exposure Ratio (SER). Delta-adjusted derivative exposure divided by liquid free float.

Wrapper–Native Growth Gap (WNG). $WNG = \Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$, where U^{native} counts native service *flows* only — fee payments, private settlement volume, proof purchases, and verified-compute consumption, each a turnover measured over the window. Collateral lockup is a stock and is published as its own companion series, never summed into U^{native} ; burns are supply arithmetic, not buyer evidence (Appendix H.14), and are likewise excluded. A persistently positive WNG means financial exposure is growing faster than monetary use — the single most compact measure of financialization outrunning adoption. Because it feeds the triple-divergence composite and Red Line 9, its aggregation is specified rather than left to the operator: a trailing four-quarter window; quantities in *unit* terms on both sides (wrapper exposure in units of X equivalent, native use in units of service transacted), so that asset-price revaluation is stripped from the comparison by construction; and the USD-denominated series published alongside for readers who want the market view. The reason for unit-denominating is a revaluation artifact, and it runs the other way from what intuition suggests: E^{wrapper} is a stock observed instantaneously and revalued with every tick, while U^{native} accumulates at transaction prices over the window, so a price *rise* — not a fall — can lift measured $\Delta \ln E^{\text{wrapper}}$ on revaluation that no usage produced. Unit terms cancel that term; the USD series exposes it. The deflation convention here is therefore *unit denomination*, which is distinct from the *work-load unit-price deflation* of Red Line 14 Condition B (a compute-price index, not the asset price): same word, different deflators, both published.

Recursive Claim Ratio (RCR). Exposure held in wrappers whose underlying is itself another wrapper or derivative, divided by total wrapper exposure — delta-adjusted on both sides, over the same window as SER. One minus RCR is the share of wrapper exposure that bottoms out in the native asset directly.

Native Use Share (NUS). The share of *observable economic activity touching the asset* that actually settles through the protocol-native monetary loop, where the universe is defined as all activity on public venues and wrappers reporting the asset (USD value of transfers and wrapper flows over the window), and the numerator is the USD value of native-loop settlement (fee-paying, proof-carrying, or corridor) over the same window. The universe is pre-declared with the metric; activity invisible to public data is excluded from both numerator and denominator and acknowledged as a coverage limit, not silently assumed zero.

24.8.3 Holder Cohorts, Vintage, and Policy Concentration

Flow elasticity is not a permanent product characteristic, and treating it as one is a modeling error. Product launch date predicts holder behavior: older funds contain embedded winners who are more likely to skim gains after a rally, whereas newly launched products may contain momentum-oriented holders with no gains to realize and therefore nothing to trim. So

$$\varepsilon_{i,t} = f(\text{product age, embedded gains, holder cohort, narrative regime, volatility})$$

Telemetry should therefore track product age, average holder cost basis where inferable, the realized/unrealized gain distribution, rolling elasticity, statistical break dates, holder-cohort turnover, and fund-flow response after drawdowns.

This matters more, not less, in an agent economy. Automated treasury systems, robo-advisers, and AI allocation agents may share common rebalancing policies, producing a machine-readable version of holder elasticity at scale. The protocol should therefore monitor concentration not only of holders but of **holder operating rules**.

Policy Concentration Ratio

The share of net asset demand governed by the largest common rebalancing templates, treasury algorithms, robo-advisers, or agent policies. Picture a thousand independent drivers, each in their own car, all running the same navigation app: ownership is distributed, but on the day the app says “take the exit,” they all take the exit. An asset can have decentralized ownership and highly centralized behavior; only the second predicts what happens on a bad day.

24.8.4 Capital Survival: AUM Is Not Investor Return

One further metric prevents a specific and common confusion. A leveraged wrapper can survive commercially while repeatedly destroying investor capital: reverse splits preserve a presentable share price, fresh inflows replenish assets, and the sponsor continues collecting fees on a product whose successive holder cohorts were harmed. Persistent AUM is therefore not evidence of successful capital preservation.

Capital Survival Ratio (CSR)

At the product level, every term must be an aggregate dollar amount — *net assets*, not per-share NAV, which would mix a per-unit price into a sum of dollar flows:

$$CSR_i = \frac{\text{current net assets} + \text{cumulative distributions} + \text{cumulative redemptions}}{\text{cumulative subscriptions}}$$

Because the stated purpose is to ask whether *successive cohorts* survived, the product-level figure is only a summary. The cohort form is the one that answers the question, for cohort c with cohort-attributable value $V_{c,t}$:

$$CSR_{c,t} = \frac{V_{c,t} + \sum_{\tau \leq t} \text{cash returned}_{c,\tau}}{\sum_{\tau \leq t} \text{capital contributed}_{c,\tau}}$$

This is a multiple on invested capital by vintage; a dollar-weighted return over the same cohort carries the same information (a cohort's CSR exceeding 1 after T quarters corresponds to a positive dollar-weighted IRR over those flows, so the two are orderings of one object, not independent evidence). A product can show CSR_i near or above 1 while late cohorts sit well below 1, which is precisely the failure the metric exists to expose. Accounting implementations vary, but the purpose does not: AUM is not investor return, product survival is not holder survival, and a wrapper can be an excellent business for its sponsor and a disaster for everyone who owned it.

24.8.5 Dealer Balance Sheets and Hedge Paths

Wrapper exposure does not necessarily arrive as direct spot purchases. Dealers may implement it with swaps, futures, baskets, or options, and counterparty capacity affects how — and whether — the exposure reaches the underlying market. Product-level exposure demand must therefore never be read as an identical quantity of spot orders.

Publish: the wrapper counterparty list; top- N dealer share; physical holdings versus swaps; swaps versus options; average derivative maturity; delta- and gamma-adjusted exposure; collateral and margin requirements; evidence of counterparty caps; and changes in hedge-instrument mix.

Dealer capacity warning. A migration from swaps toward options can indicate that dealer balance-sheet capacity is binding. When dealers are short calls, rising spot and rising volatility can require additional underlying purchases, adding a gamma feedback loop on top of the mechanical one. This is an empirical warning to watch for, not a universal law to assume.

24.8.6 Convenience-Yield Telemetry: Pricing the Holder-Side Service Flow

VERIFYPRICE, VerifyReach, and VERIFYSETTLE measure the stack; the wrapper vector measures the realization plane; the Native Monetary Buyer Map measures who is holding and how. None of them prices the thing the monetary claim actually rests on. [Section 11.10.8](#) relocates the premium to a **regime-contingent convenience yield** — a holder-side service flow with no cash flows to discount — and then concedes it is unsized. The concession is hon-

est but incomplete: unsized does not mean unmeasurable, and an unsized quantity with no instrument is indistinguishable from an article of faith. Storage theory extracts convenience yield from observable prices rather than asserting it (the forward-curve slope net of carry); the crypto-native analog is the lending and derivatives complex around the asset. The thesis therefore publishes a fifth measurement family whose job is to make the premium a number rather than a hope.

Instruments. For each venue publishing a rate, report:

Base-asset lending rates. The market-clearing rate for borrowing the base asset itself, segmented by venue, term, and counterparty class. A borrower pays this rate to be short the asset; what they are paying *for* is the units. When holders of the bearer service lend out their units, they are selling the service flow for the term of the loan — so the lending rate is the direct market price of temporary access to bearability, the closest available analog to a lease rate.

Forward and perpetual basis. Forward price minus spot, and perpetual funding rate, net of the reference interest rate. Persistent positive basis that exceeds carry costs indicates demand for *deferred or levered exposure* outrunning demand to hold the units — demand for the price story rather than the bearer service. Persistent negative basis indicates the reverse: entities that need the units (operators meeting collateral calls, shorts delivering into settlement) bidding against the float.

Wrapper basis. The premium or discount of a custody-restricted wrapper over the value of its native holdings. A persistent premium is the price of *not* holding natively: legal wrapper, custody, and access convenience purchased in place of self-custody, agency, and exit. It is the complement of the lending rate — both price the friction between native holding and representative holding — and under [Corollary 11.1](#) a large and rising wrapper premium with stagnant native use is financialization, not monetary adoption.

What the instruments read. Convenience yield is a *residual claim* on holding, extracted from prices that also contain discount-rate and risk information, so each reading requires the same care as ε in [Appendix H](#):

- **Lending rate decomposition.** Borrow demand has three sources: directional short interest, market-making inventory financing, and operational short squeeze against locked float. Segment by borrower class where observable; an operational squeeze (collateral calls meeting immobile supply) is a liquidity event, not a premium reading.
- **Regime conditioning.** The thesis's claim is specifically *regime-contingent*: the yield should rise in exactly the states where substitutes weaken. The instrument earns its place only if the rate series is published against the macro-state model of [Section 27.2](#) and the pressure index R of [Section 11.14.1](#), so that convenience yield = $g(R)$ becomes an estimable function

rather than a slogan. A yield flat across regime states is the null result and must be reported as such.

- **Both tails checked.** A convenience yield that appears only in benign states — holders renting out units for yield-farming while nobody needs the bearer service — is a fee-like cash flow, not the regime-contingent service flow, and belongs in the DCF bucket with the rest of [Section 11.10.7](#).

The null result is publishable. The instrument is falsification machinery, not advocacy. If, after the stack delivers measurable service, the lending and basis series show no regime-contingent premium — holders charge nothing extra to part with units exactly when substitutes fail — then the holder-side flow is priced at zero and the store-of-value claim has lost its mechanism in public. That outcome belongs on the Value Capture Board beside fee coverage, and it is a more direct test than any red line currently operating on it: the red lines test whether the *chain* broke; this tests whether the *premium* exists.

24.8.7 Buyer Quality and Price Provenance

Wrapper vectors answer *which product* is transmitting exposure. They do not answer *who is absorbing the risk*, with whose capital, under what mandate, at what leverage, against which liability, through which benchmark, and with what capacity to sit through a mark. That is the sovereign-market question of [Section 3.1.2](#), applied to the native asset's own realization plane. Publish, as a checklist rather than as a multiplicative index:

- **Persistence (*P*).** Does the bid recur because contributions, a mandate, or a liability continue, or because a trade has not yet finished?
- **Horizon (*H*).** Is the holder matching a multi-decade promise, a quarterly risk limit, or a daily reset?
- **Liability match (*M*).** Is the purchase offsetting an enduring obligation, or following a weight?
- **Loss-bearing capacity (*B*).** Can the holder absorb marks without forced sale, margin call, or redemption?
- **Countercyclicity (*C*).** Does incremental buying increase when prices fall, or does the governing rule absorb less risk precisely then?

A high-persistence, low-match, low-countercyclicity buyer is the automatic index flow of [Section 3.1.2](#): it keeps arriving and still fails as an anchor. A leveraged relative-value fund can score the opposite way on persistence and the same way on loss-bearing capacity. Neither score is a number to multiply into a headline Q_A . The fields exist so that “there is demand” cannot be reported without saying *what kind*.

Price then requires a provenance model, not only a feed: who bought, with whose capital, under what mandate, at what leverage, with what redemption terms, against which liability,

through which benchmark, and with what capacity to survive losses. A policy that treats the resulting price as an external referendum on credibility is reading an instrument the state and the financial industry themselves constructed. That is the epistemic loop [Section 3.1.1](#) already warns about, now stated as a VerifyFlow requirement: **no buyer-quality ledger, no inference from the long yield—and no inference from the native asset’s price either.**

24.8.8 Native Monetary Buyer Map

Buyer quality belongs inside the monetary chain, not only in the external price model. Publish a Native Monetary Buyer Map that reconciles protocol flows and holder cohorts without pretending that pseudonymous addresses are known persons:

Just-in-time fee acquisition. Share of fee-paying units acquired within pre-declared windows before expenditure, versus fees paid from retained balances.

Operator sell-through and inventory. Rewards and fee receipts sold, retained, hedged, or pledged by operators, by cohort and interval.

Burns and net issuance. Gross burns, gross issuance, and net issuance separately; a burn is supply arithmetic, not evidence of a buyer.

Self-custodied reserve demand. Net accumulation in addresses or proofs-of-control classified under published, conservative reserve criteria, excluding custodial and wrapper inventories.

Wrapper demand. Custodial, fund, treasury-company, and synthetic exposure reported apart from native reserve holdings.

Leverage. Borrowed funding, derivatives, rehypothecation, liquidation sensitivity, and margin terms where observable.

Holding period. Cohort survival and realized holding-duration distributions, not a single average that mixes fees with reserves.

Loss-bearing capacity. Evidence that the funding source can absorb marks without margin call, redemption, or forced sale.

Countercyclical accumulation. Whether native self-custodied reserve cohorts add units during drawdowns and regime-pressure episodes rather than merely after recoveries.

The map must reconcile identities only to the degree evidence permits. Unknown holdings remain unknown; they are not assigned to the most favorable category. Gross native volume passes Premise 7 only when the retained, self-custodied, unlevered or transparently financed, loss-bearing, countercyclical cohort is material and persistent.

The map also carries a representation discipline that every published aggregate in this chapter inherits: any aggregate over households or holders must disclose its weighting rule. The weighting rule is a moral choice made visible. Dollar-weighted and person-weighted summaries of the same underlying observations can move in opposite directions — and neither is the false one. The arithmetic is exact. The gap between the two means is the covariance of

spending weight and experience over mean spending weight, $\bar{x}_{\text{dollar}} - \bar{x}_{\text{person}} = \text{Cov}(C_i, x_i) / \bar{C}$, so a positive covariance between spending weight and experience lifts the dollar mean while the person median falls. That is the mechanism behind the “two economies” reading of the same dataset (“The Vibeceession Was Real”; Sources). Board aggregates therefore publish their weights (units transacted, addresses, identities, delta-equivalent exposure) alongside the number, and where a dollar-weighted and a person-weighted reading of the same series diverge persistently, both are shown rather than the more favorable one. A verified aggregate that describes almost no actual participant is not a measurement failure; it is a choice of weights, and the choice is now visible.

24.9 The Public Boards

Metrics surface as twelve public boards or maps. The original five have grown to twelve as the thesis has extended from pure verification economics into physical capacity, political economy, dependency, holder quality, and market structure:

Proof & Compute Board. Summarizes VERIFYPRICE per workload, SLA attainment, queue depths, and failure modes.

Settlement & Privacy Board. Reports swap success, refund safety by corridor, time-to-finality, and anonymity-set health.

Neutrality & Admission Board. Tracks time-to-first-proof for new entrants, top-N share, geo/ASN distribution, and house share.

Layer 0 Capacity Board. Extends hardware-profile telemetry with Facility Energy Receipts and Facility Capacity Receipts (Section 15.7): grid node, interconnection class, outage history, backup duration, fuel mix, cooling dependency, and Physical VerifyPrice (Section 20.8). Reports both Sovereign Optionality and scenario-specific Delivered Verified Capacity — how much of the grid still lights up when the weather turns — plus active minimum cuts, substitution latency, and dependency groups (Section 15.9). Tracks whether the compute/proof layer is physically resilient or just another hyperscaler/cloud bottleneck, including the **Homestead Ratio** (Section 5.6).

Economic Coverage Board. Charts fee-coverage ratio, physical VERIFYPRICE against SLOs, and demand curves indexed to policy events.

Value Capture Board. Tracks fee coverage, native-asset burn rates, collateral lockup volumes, issuance vs. capacity growth, workload mix (budgeted vs. speculative), and bypass-channel indicators (operator off-ramp rates, stablecoin-denominated fee share). This board makes the **Value Capture Lemma** from Chapter 7 auditable: if value capture is failing, it shows here before it shows in price. It also carries the **triple divergence** composite — persistently positive WNG, flat or falling regime-conditioned convenience yield, and falling Homestead Ratio against improving closed-stack service metrics — which is the observable signature of the joint

failure mode named in [Section 30.9](#).

Wrapper Dominance Board. Tracks whether the asset is becoming a financial wrapper rather than money: the **Wrapper Dominance Ratio** ([Section 11.13](#)) alongside its component series (ETF/wrapper AUM, custodied balances, exchange balances, wrapped supply, native fee share, non-custodial settlement volume, shielded settlement volume, collateral lockups). A rising WDR while native usage stagnates is surfaced here before it appears anywhere else.

Agency Preservation Board. Tracks non-custodial usage share, selective-disclosure scope, dossier-reconstitution risk, and forced-disclosure incidence ([Section 25.1.2](#)), operationalizing the ninth SoV requirement ([Item 9](#)). This board is optional in early phases but becomes load-bearing once the protocol has material institutional and administrative-state interaction.

Market Realization & Wrapper Board. Reports the VerifyFlow vectors ([Section 24.8](#)): wrapper inventory by class, assets and delta-equivalent exposure, creations and redemptions, holder flow elasticity, net mechanical gain, Wrapper Recycling Ratio, constituent concentration, market-depth ratios, hedge-instrument mix, counterparty concentration, the Wrapper–Native Growth Gap, custodial versus self-custody shares, recursive wrapper exposure, cross-jurisdiction exposure, holder-cohort statistics, and the convenience-yield telemetry of [Section 24.8.6](#) (lending rates, forward and perpetual basis, wrapper basis, published against the regime-pressure index).

Native Monetary Buyer Map. Separates just-in-time fee acquisition, operator sell-through and inventory, burns and net issuance, self-custodied reserve demand, wrapper demand, leverage, holding period, loss-bearing capacity, and countercyclical accumulation ([Section 24.8.8](#)). It asks who warehouses native monetary risk, not merely whether demand exists.

Common-Cause Dependency Board. Maps dependencies crossing nominally independent facilities or jurisdictions: hardware and accelerator vendors, firmware and signing systems, energy and fuel markets, cables and network routes, identity, custody, market making, banking, legal recognition, cooling-water basins, and software distribution. For each dependency publish affected capacity, substitutable alternatives, substitution latency, correlated stress scenarios, and the DVC loss if the group fails. The correlated physical-failure framing is adapted selectively from Doomberg’s *A Minor Detail*, treated as scenario commentary rather than forecast.

Sponsor Dependency Board. Reports subsidy concentration, politically contingent demand, cloud-credit dependence, dominant-customer exposure, state or foundation veto power, operator exit capacity, sponsor-withdrawal sensitivity, conflicting terminal objectives, capacity removable on short notice, and spillover to neutral third parties. Subsidized operators and sponsors are separate strategic actors even when their current incentives align. The principal–agent framing is adapted selectively from Doomberg’s *Over Target*; inferred intent and tactical escalation claims are excluded.

Do not merge the last two economic boards. The Value Capture Board and the Market

Realization & Wrapper Board answer different questions, and collapsing them destroys the distinction the thesis depends on:

- **Value Capture Board:** *Is service value accruing to the native asset?*
- **Market Realization & Wrapper Board:** *What is currently setting its price?*

The Value Capture Board stays focused on protocol economics. The Native Monetary Buyer Map asks whether a monetary-risk warehouse exists. VerifyFlow stays focused on market realization. A reader needs all three to distinguish service accrual, monetary adoption, and financialization.

These boards are not marketing pages; they are part of the protocol's public interface.

24.10 Detecting Drift

Drift is the failure mode that never announces itself. Nothing breaks; everything erodes — a few milliseconds added to p95 here, a few points of share consolidated there — and each reading is individually explicable, which is precisely why an adversary (or mere entropy) prefers it to an attack. Telemetry exists to catch the trend while it is still a trend:

Verification cost creep. VERIFYPRICE shows p95 verify time creeping up. Response: R&D lab proposes fixes; governance ratifies changes.

Verifier concentration. 80–90% of verifications performed by one hardware profile. Response: adjust rewards to subsidize diverse profiles.

Reachability degradation. VerifyReach shows a major country blocking protocol ports. Response: shift transports, promote alt routes. You can dam a river; you cannot dam the rain — the response is always to add paths faster than the adversary can close them.

Corridor fragility. VERIFYSETTLE shows corridor success dropping. Response: re-weight corridors, onboard new LPs.

24.11 Making Neutrality and Repression-Resilience Falsifiable

Neutrality and repression-resilience are usually sold as moral properties. In this thesis, they are **falsifiable hypotheses**:

Neutrality hypothesis: “No actor or group can reliably censor, privilege, or front-run specific flows across Layers 1–5 without being detected and without other actors being able to route around them.”

Repression-resilience hypothesis: “Under YCC, capital controls, DPI, and blacklists, the triad's core capacities retain sufficient DVC and the base asset's self-custodied holder constituency retains sufficient agency to remain a monetary candidate.”

Telemetry makes these hypotheses testable:

- If VerifyReach shows certain flows consistently blocked and no alternatives emerge, neu-

trality is falsified.

- If VerifySettle shows privacy corridors collapsing under regulation, repression-resilience is falsified.
- If VerifyPrice shows verification cost pushed beyond reach, “anyone can verify” is falsified.
A system that cannot tolerate bad news cannot be money; it can only be marketing.

24.12 Policy-Attack Stress Harness

A monetary stack designed for repression should publish the exact conditions under which it considers itself passing or failing. For each attack scenario, we define tests and pass criteria:

Yield-curve control & negative real yields:

- **Test:** Over 24–36 months of -300 to -500 bps real yields, fee+burn share of security budget is stable or rising.
- **Pass:** Tokens behave like claims on useful capacity, not synthetic bonds.

On-/off-ramp throttling:

- **Test:** Non-custodial corridors maintain $\geq 95\%$ success and 100% refund safety; anonymity sets remain above thresholds.
- **Pass:** Users can move capital via non-custodial rails even when custodial ramps are hostile.

Provenance mandates:

- **Test:** Canonical provenance workloads have healthy VerifyPrice; regulators can run `verify(receipt)` with commodity tooling.
- **Pass:** The triad is the cheapest way to comply with provenance mandates.

Hardware backdoor mandates:

- **Test:** A non-trivial share of receipts are tied to open or sampled hardware profiles, not a single opaque TEE.
- **Pass:** Hardware mandates become one option among many, not a kill switch.

When the stack claims “we can survive policy attacks,” the proof is not a blog post; it is a corpus of **stress-harness runs with metrics and receipts**.

24.12.1 Market-Realization Stress Tests

The harness above stresses repression and availability. The Market Realization Plane needs its own scenarios, because the compositional adversary (Section 5.2) does not respond to any of the tests above. Each scenario is run against the VerifyFlow vectors of Section 24.8.

A. Wrapper inflow cannon. A large spot or thematic fund launches; creations arrive independent of daily returns; holdings are concentrated; native liquidity is shallow. **Measure:** Q^{RD}/V , constituent concentration, price impact, native-use response, WNG. **Question:** does the bid convert into native usage, or only into price?

B. Elasticity regime break. A $2\times$ or $3\times$ wrapper grows rapidly and holder elasticity drifts toward zero, so recycling declines. **Measure:** κ , Wrapper Recycling Ratio, volatility feedback, reversal symmetry. **Question:** has a product quietly become a momentum amplifier?

C. Creation halt. Wrapper inflows stop while protocol-native usage stays constant, removing passive support. **Measure:** price drawdown, native fee coverage, governance response. **Pass:** native fee coverage remains sound *and* governance does not respond by inflating issuance or subsidizing price.

D. Overnight gap or market halt. The underlying gaps beyond a daily-reset wrapper's viable range before it can rebalance. **Measure:** wrapper solvency, dealer collateral calls, liquidation paths, cross-venue settlement. **Pass:** native protocol operations continue unaffected; wrapper failure stays contained in the Market Realization Plane.

E. Swap-capacity exhaustion. Dealers reduce swap exposure and funds migrate to options. **Measure:** implied-volatility demand, gamma hedging, dealer concentration, collateral requirements, spot-volatility correlation.

F. Price up, protocol flat. Price triples and wrapper AUM grows while native fees, settlement, proofs, and compute usage do not. **Expected result:** market success, monetary-adoption failure, and an *automatic* dashboard warning. A thesis that cannot raise an alarm during a rally is not falsifiable.

Chapter 25

Legal, Policy, and Jurisdictional Posture

The adversary here does not send exploits; it sends letters. Layer 6 sits where protocol meets law and politics, and the encounter is asymmetric: code can be amended in a release cycle, but sanctions regimes, travel rules, and court precedents move on their own clock and do not consult the changelog. If the triad is to be sustainable, it must support **lawful privacy**, defend itself against bans without being defined by them, and avoid capture by any single jurisdiction. That means designing for the law as it is — including where it is hostile — rather than for the law as the cryptography would prefer it.

25.1 Lawful Privacy as Protocol Design

“Lawful privacy” is often a euphemism. Here it means something precise:

- **By default**, settlement and compute flows are private.
- **By design**, participants can opt into **selective disclosure** via viewing keys and receipts.

Constitutional Constraint: Policy = Predicates

Policy compliance is defined as **ZK-checkable predicates** over credentials and receipts, plus **selectively disclosed audit windows** via viewing keys.

It is **not** defined as global graph inspection or universal traceability. Any policy requiring ubiquitous transaction tracing is treated as **incompatible with the monetary design**.

25.1.1 What Lawful Privacy Is Not

Lawful privacy is not a promise that users can ignore legal obligations. It is a design posture in which the protocol minimizes unnecessary disclosure while allowing users to produce scoped, auditable evidence when they choose or are legally required to do so. Specifically:

- Lawful privacy does **not** guarantee immunity from subpoena, tax reporting, sanctions enforcement, or employment obligations.

- It does **not** mean that regulators will accept predicate-only compliance in all jurisdictions. Some jurisdictions will reject it.
- It does **not** mean that users can evade capital controls or hide from governments. The protocol provides *technical* confidentiality, not *legal* immunity.
- It **does** mean that the protocol itself does not require surveillance infrastructure, and that disclosure is always at the holder’s discretion rather than built into the protocol.

The legal posture is **defensive, not antagonistic**: the thesis designs for resilience under legal heterogeneity, not for confrontation with law.

25.1.2 Social Coercion Metrics: When “Optional” Becomes Mandatory

The protocol already acknowledges that lawful privacy is a technical guarantee, not a social guarantee (Section 5.6): it cannot prevent coercion to disclose keys, but it can raise the cost of mass surveillance, preserve optionality, and make coercion visible through telemetry. The remaining question is *when* optional disclosure becomes mandatory in practice.

Forced Disclosure Incidence. Reported cases where employers, banks, landlords, visa processes, exchanges, or benefits agencies require viewing keys.

Selective Disclosure Scope. Average breadth of information disclosed per audit event.

Dossier Reconstitution Risk. Probability that repeated selective disclosures can reconstruct the user’s full graph.

Non-Custodial Practicality Index. Share of flows that can be completed without a KYC choke-point at equivalent service quality.

Lawful privacy fails socially before it fails cryptographically. If viewing keys become routine prerequisites for employment, banking, housing, travel, or benefits, optional disclosure becomes administrative disclosure—the social-layer expression of the Administrative Repression pattern introduced in Section 5.2. The protocol cannot stop all coercion, but it can measure whether coercion is becoming systemic. These four metrics feed the Agency Preservation Board (Section 24.9).

Requirement	How Satisfied (Protocol)	What Current Law Actually Requires
AML screening	ZK proof of “sender $\in$ cleared-set”	Originator identity data (Travel Rule)
Jurisdictional limits	ZK range proof of amount	Exact amount above reporting thresholds
Tax audit	Time-bounded viewing key	Records sufficient to reconstruct transactions
Counterparty verification	ZK set-membership proof	Named-party identification

Table 25.1: Lawful privacy: what the protocol can satisfy versus what current law demands. The right-hand column is not a design boundary the protocol can hold by engineering; it is the law as written, and predicate compliance is *unlawful* under it in most major jurisdictions until the law changes.

Read the table honestly. An earlier draft of this table headed its right column “Cannot Require,” implying the protocol holds a boundary that regulators must respect. That framing is wrong and has been corrected. The Travel Rule (FinCEN; FATF Recommendation 16; the EU Transfer of Funds Regulation) obliges obligated entities to *transmit and retain actual originator and beneficiary identity data* — name, account, amount — not set-membership proofs. A regulated institution cannot lawfully satisfy that obligation with predicates, no matter how sound the cryptography. The accurate statements are:

- **Within the current law, predicate compliance is an engineering option that is not yet a legal option.** It requires legislative or rule-level change in every major jurisdiction before a regulated counterparty can rely on it. The protocol’s contribution is to make that change *expressible*: when a jurisdiction moves, the machinery exists to comply with far less surveillance than universal tracing would need.
- **Outside regulated perimeters, the predicates are simply private.** P2P settlement between non-obligated parties is not a Travel Rule event anywhere today. The design buys lawful headroom at the edges while the institutional perimeter remains a policy question, not a protocol one.
- **Who curates the cleared set is a hard, unresolved question.** A “sender $\in$ cleared-set” proof requires an oracle maintaining the set: a sanctions-list curator inside a privacy protocol is simultaneously a chokepoint, a legal attack surface, and a party with liability. The thesis does not have an answer and declines to pretend otherwise.

25.2 Securities Law: The Value-Capture Mechanism Is Its Own Legal Risk

The strongest legal threat to the asset is not a ban. It is classification, and the threat comes from inside the thesis rather than outside it. The [Value Capture Lemma](#) is, in its design commitments, a fact pattern regulators have a name for: fees routed to holders by protocol decision (burns as pro-rata value return), returns derived from the efforts of operators and developers, an expectation of profit from holding, and governance-adjustable fee splits. Under the Howey framework — investment of money, in a common enterprise, with a reasonable expectation of profit derived from the efforts of others — each element is arguably present, and the stronger the value-capture design *as an economic argument*, the stronger the securities case *as a legal one*. The irony is structural and must be stated plainly: **the thesis’s best economics is affirmative evidence for the classification that would end institutional holdability in the jurisdictions that matter most.**

- **A burn is a distribution.** Economically the thesis concedes it is a buyback ([Section 11.10.7](#)); legally a buyback is a managerial act returning value to holders. “The burn is not a dividend” is an engineering statement, not a legal one.

- **Governance-adjustable fee routing is discretionary management.** The reference designs let token governance alter fee splits — the 70/20/10 default, treasury share, burn fraction. Discretion over the return to holders is the “efforts of others” prong wearing a multisig.
- **The checklist cannot treat classification as exogenous.** Chapter 26 asks whether the asset’s legal classification is clear enough for treasury or fund mandates, as though clarity were weather. Classification is partly a *consequence of design choices made in this thesis*, and the choices that maximize value capture are the choices that maximize Howey exposure.

This is a genuine design tension, not a footnote. The available escapes move in opposite directions, and the thesis should name the corridor rather than drift through it. Removing holder-returning mechanics (no burns, fee destruction only at point-of-sale, purely consumptive fees) weakens Conditions 2–3 of the lemma and moves the asset toward a pure service medium — legally cleaner, monetically thinner. Keeping them maximizes the cash-flow claim and the securities exposure together. A third path — issuance and fee policy fixed *ex ante* in a constitution beyond governance reach, with no discretionary re-routing ever possible — is the only version of value capture that does not depend on “efforts” anyone could alter, and it is also the version closest to what a store-of-value constitution should look like anyway (Section 23.6). Red Line 7 tracks whether lawful participation remains possible; the securities question determines whether Red Line 7 can be satisfied at all in the largest capital pools.

This section is analysis, not advice. Jurisdictions differ, enforcement posture shifts, and no statement here is legal advice. The point of writing it down is that a thesis which publishes fifteen red lines but leaves its largest legal vulnerability implicit would be failing its own standard.

25.3 Defense-in-Depth Against Bans and Sanctions

Regimes may attempt to ban privacy assets, sanction addresses, mandate KYC at all edges, or classify triad services as “unlicensed financial institutions.”

The live precedent for this section is not a statute but a sanctions action: the designation of Tornado Cash by the US Treasury’s OFAC in 2022 — and the subsequent prosecution of its developers — established that a non-custodial, autonomously operating privacy protocol *as such* can be treated as a sanctionable entity, and that writing and publishing mixing code can expose its authors to liability. The action was partially reversed on appeal in 2024 (the Fifth Circuit held immutable smart contracts are not “property”), but the operating lesson survives the reversal: the threat vector for privacy infrastructure has not primarily been domestic statutes; it has been secondary-sanctions reach through USD intermediaries, exchange delisting under compliance burden, and personal liability for developers. Three consequences for this stack:

- **Developer and operator liability is a first-class risk surface**, not a footnote to Layer 5’s

corridor design. Anonymous or distributed development, jurisdictional separation of maintainers from operators, and legal entity structure (Section 25.4) are not conveniences; they are the difference between a delistable product and a prosecution target.

- **Exchange delisting is the working enforcement mechanism, not the ban itself.** Monero and Zcash were not outlawed in most jurisdictions; they were made practically inaccessible by travel-rule compliance burden at regulated venues. Corridor design must assume the fiat on-ramp is the chokepoint that closes first (Section 11.12.1 engages the substitute channel this creates).
- **Non-custodial is a legal argument, not a shield.** The strongest available defense is precisely the non-custodial architecture this thesis specifies — no controller, no admin keys, no ability to freeze — because it is the fact pattern that defeated the Tornado Cash designation in court. The design should maximize those properties deliberately, because they are load-bearing in litigation, not only in threat models.

Layer 6 cannot prevent law from existing, but it can:

- Avoid giving any single jurisdiction a **kill switch**.
- Maintain **jurisdictional diversity** in hardware profiles, proof factories, corridor LPs, and foundation incorporation.
- Provide **fallback modes**: local-first clients, mesh/satellite relays, minimal CLI modes that maintain basic functionality under degraded conditions.

The legal posture is **resilience under legal heterogeneity**: the stack does not advertise itself as a tool to evade law, but designs for continued operation across a diverse legal landscape where some jurisdictions are friendly, some hostile, and most ambiguous.

25.4 Labs, Foundations, and Neutral Router Commitments

Key commitments:

Neutral router charters. Router operators commit to content-agnostic routing. Deviation is punished in-protocol and reputationally.

Foundation / lab structure. Legally separate entities with public charters that define roles and explicitly renounce certain powers.

Multi-jurisdiction footprint. Incorporations and key staff spread across legal regimes to avoid single-point capture.

Chapter 26

Operator & Investor Checklist

The last piece of Layer 6 is a practical checklist: how builders, operators, and allocators decide whether this stack qualifies as a next-generation store of value. It is deliberately unglamorous, and that is the point. Everything before this chapter argued the thesis at full strength; a checklist is where you stop believing the argument and start checking it. Read it as the stack's own answer to "what would change your mind?" — each question below has a yes that strengthens the claim and a no that weakens or falsifies it.

26.1 Triad Supply: Privacy, Proofs, Compute

For each leg of the triad:

Privacy:

- Are there live, non-custodial privacy corridors with measurable anonymity sets?
- Are shielded pools growing in depth and churn?
- Are lawful-privacy patterns (viewing keys + receipts) actually used?

Proofs:

- Is there robust supply of proofs for canonical workloads?
- Are VERIFYPRICE SLOs met consistently?
- Is proof production diversified across hardware profiles and geographies?

Compute:

- Are meaningful amounts of useful compute being proven?
- Is there a liquid market for Work Credits?
- Are triad workloads tied to real-world demand?

26.2 Value Capture & Anti-Bypass

The most important questions for evaluating whether the SoV thesis holds:

- Can users buy equivalent Privacy, Proofs, or Compute service **without the native asset**?

- Are fees **actually paid in the asset**, or do most 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) that captures most of the demand without flowing through the asset?
- Does governance have the power to **redirect value away from holders** (e.g., by eliminating burns, inflating supply, or changing fee routing)?
- Is the **Wrapper Dominance Ratio** (Section 11.13) 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 (Corollary 11.1)?

If the answer to any of these is “yes,” the system may be useful infrastructure but the SoV thesis is weakened or falsified.

26.3 Institutional Holdability

- Is **custody support** available from regulated custodians?
- Is the **legal classification** clear enough for treasury or fund mandates?
- Is the asset **auditable** enough for institutional reporting requirements?
- Is **liquidity** deep enough for institutional-scale positions?
- Is there a **valuation policy** that allocators can apply (fee-flow models, comparable frameworks)?

Holdability is a virtue with a bill attached (Section 5.6). Every question above should therefore be paired with its market-realization counterweight:

- Does wrapper growth **seed** native usage, or **substitute** for it? Is the Wrapper–Native Growth Gap (Section 24.8) persistently positive?
- What share of supply sits under custodial control (CCR), and how concentrated is that custody?
- Is delta-adjusted synthetic exposure (SER) large relative to liquid free float?
- Do daily-reset leveraged products exist on the asset, and what is their net mechanical gain κ and Wrapper Recycling Ratio?
- Is exposure concentrated in a small number of dealers or counterparties, and has the hedge-instrument mix shifted from physical to swaps to options?
- Is demand governed by a small number of common rebalancing policies (Policy Concentration Ratio)?

26.4 Stack Health: Layers 0–6

Layer 0. Documented hardware profiles? Open hardware participation? Visible incident reports? Facility Capacity Receipts available and independently verifiable (Section 15.7)? Physical VerifyPrice within bounds (Section 20.8)?

Layer 1. VerifyReach shows resilient connectivity? Alternative transports exist?

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

Layer 3. Pseudonymous credentials integrated? Actors can prove rights without doxxing? Non-custodial usage share and dossier-minimization score visible (Item 9)?

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 processes? Recent incident reports?

Market Realization Plane. Is VerifyFlow published at all (Section 24.8)? Are wrapper inventory, creations/redemptions, flow elasticity, concentration, and the Wrapper–Native Growth Gap visible? Is the market-realization stress harness (Section 24.12.1) run and reported? Remember that this plane sits *around* the stack, not inside it: its failure does not break the protocol, but its opacity breaks price interpretation.

26.5 Telemetry Honesty

Telemetry itself can be gamed. Checklist for honesty:

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

26.6 Red Flags and Failure Patterns

Finally, some patterns that should trigger skepticism:

Closed hardware monoculture. One TEE, one vendor, no sampling, no open profiles → Layer-0 cliff.

Opaque bridges. “Magic multisig” bridges with no proofs, unclear jurisdiction, and no incident history → not settlement rails, just custodial risk.

Foundation fiat. Major parameter or policy changes via blog post, with no on-chain trace, no telemetry, and no incident report → governance capture.

Zombie corridors. Privacy rails that haven’t moved meaningful volume in months but still appear in marketing. Anonymity sets that are effectively dead.

Proof theater. Lots of “ZK” branding but no public VERIFYPRICE metrics, no canonical workloads, no commodity verifiers.

Governance theater. Token votes with single-digit participation deciding fundamental parameters; no constraints; no SLOs.

Price as proof. Marketing that cites market capitalization, wrapper AUM, product launches, or trailing returns as evidence that the monetary thesis is working, while native fee, burn, collateral, and settlement series are absent or flat → financialization narrated as adoption (Corollary 11.1).

Any one of these is survivable in the short run; taken together, they say:

“This is not a triad-backed money system; it is a platform dressed as one.”

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.

Part V closes the governance loop:

- It treats SLOs and dashboards as the **constitution**.
- It makes **Verify* telemetry** the lifeblood of neutrality and repression-resilience.
- It frames law and politics as **constraints to be engineered around**, not as deities to be appeased.
- It gives builders and allocators a **practical checklist** for deciding whether a system’s “store of value” claim is mathematically grounded or just well-typeset.

With Layer 6 in place, the stack has all three angles: Money (triad as SoV), Stack (Layers 0–6), and Telemetry (Verify* keeping it honest). Part VI can now focus on dynamics (adoption curves, risk, and implementation sketches): how this thing actually gets built, traded, attacked, and, if it works, quietly becomes part of what the world calls “money.”

Part VI

Dynamics, Risk & Implementation

Parts I–V defined the triad as distinct services that may support a conditional base-asset monetary candidate, built a seven-layer stack to supply them, and wired in governance and telemetry so deliverability, holder quality, neutrality, and repression-resilience are falsifiable. Part VI asks the unromantic question:

What happens when this stack hits the real world? Who uses it first, how does it fail, and under what conditions does the separate base asset remain a credible monetary candidate?

Chapter 27

Adoption Curve & Ecosystem Dynamics

The stack is not adopted by “the market” in one step. It’s pulled in by different constituencies, each with their own pain.

27.1 The Transition Window: Hard Assets as Bridge

Before the triad stack is operational, rational actors face a dilemma: soft guarantees are visibly failing, but the new infrastructure isn’t yet robust enough to absorb serious capital. Gold and silver—time-tested, non-digital, seizure-resistant in physical form—remain sensible bridge assets during this window.

This thesis does not argue against holding them. It argues that once VerifyPrice, VerifyReach, VerifySettle, and DVC are live and healthy, the triad stack may offer distinctive digital services:

- **Verifiability:** Gold’s purity requires assay; proofs verify in milliseconds on a laptop.
- **Programmability:** Gold cannot enforce conditional logic; typed Work Credits can encode service SLAs, escrow, and policy without becoming money.
- **Portability under surveillance:** Physical gold is detectable at borders; shielded settlements are not.
- **Divisibility and settlement speed:** Gold settles in days with intermediaries; privacy corridors settle in minutes without custody.
- **Service accrual:** Triad usage can route fees and burns to the associated base asset, producing a cash-flow claim and floor rather than proving monetary premium.

The question is not *whether* to hold hard assets, but *when* the new stack becomes robust enough to migrate. The adoption curve below provides the metrics: once Phase I gate conditions are met (VerifyPrice dashboards live, corridors refund-safe, receipt volume non-trivial), the rational rebalancing can begin. Until then, gold and silver remain what they have always been—stores of value that do not require trust in any single institution, even if they lack the programmability and verification properties of the triad.

Gold is the bridge; the triad is a candidate destination if telemetry and adoption gates are met. The bridge is still load-bearing while the destination is under construction.

27.1.1 Portfolio Migration: From Bridge Assets to Triad Assets

Migration should be telemetry-gated, not narrative-gated. Gold, silver, Bitcoin, front-end liquidity, and physical infrastructure remain the bridge while the destination is under construction. Triad exposure should scale only as the stack clears observable gates:

Gate 0 — Bridge regime. Hold old-world hard assets and liquidity; triad assets are venture/research exposure.

Gate 1 — Verification. VERIFYPRICE dashboards are live for canonical workloads; p95 verification cost remains within SLOs across multiple quarters.

Gate 2 — Settlement. VERIFYSETTLE shows non-custodial corridors maintaining high success and refund safety under stress.

Gate 3 — Reachability. VERIFYREACH shows access across jurisdictions, ASNs, and censorship regimes.

Gate 4 — Value capture. Native-asset fee share, burns, collateral lockups, and non-bypassability are material and publicly visible.

Gate 5 — Institutional holdability. Legal classification, custody, accounting, liquidity, and operational controls allow non-speculative treasury allocation.

Gate 6 — Repression stress. Under a real or simulated repression episode, the stack remains usable and fee/burn coverage is stable or rising.

Before Gate 1, the asset is research exposure. After Gate 3, it is infrastructure exposure. After Gate 4 and Gate 5, it becomes a credible store-of-value candidate. After Gate 6, it begins to justify reserve-asset language.

This does not mean hard assets become obsolete once the triad matures. Gold requires no network, no software, no telemetry, no cryptographic agility, and no functioning power grid; it remains rational non-digital redundancy even in a mature-triad world, not merely a placeholder for it.

27.2 Macro Policy State Model

AfterFiat is a structural architecture, not a tactical Treasury short. The same sovereign system can pass through different states, and policy repair can extend incumbent arrangements for years. The model therefore specifies states and observables, not a dated sequence:

The market-repair distinction is adapted selectively from Michael Green's commentary in *What the Treasury Needs*; the observable repair sequence and its falsifiers from *Sometimes, You Get What You Need*; and the duration-buyer tripwire from *Anchors Aweigh, My Boys*. They

supply mechanisms and scenarios, not adopted quantities or tactical forecasts.

State 1 — Market repair succeeds. Long yields are driven principally by term premium and impaired duration absorption while long-dated inflation compensation remains contained. Buybacks, maturity changes, regulatory repair, or lower front-end rates restore market function; real rates fall, mortgage duration can shorten, and private substitutes remain credible. The immediate repression premium for an open bearer stack may narrow even while the secular agency case remains.

State 2 — Physical inflation blocks repair. Energy, refining, gas, power, shipping, cooling, transformer, or industrial bottlenecks lift long-dated inflation compensation. The financial system may need easier policy while the physical system argues for restraint. Layer 0 costs rise at the same time demand for neutral assets may rise. Long-dated breakevens and inflation swaps are the tripwire distinguishing this state from a pure term-premium repair.

State 3 — Market repair fails and repression follows. Liquidity support, issuance changes, or easing do not restore a durable marginal buyer, and the state shifts from market repair toward captive demand, preferential regulation, negative real returns, capital controls, or restricted exit. This is financial repression, not merely intervention.

State 4 — Competent closed-stack stabilization. The state successfully coordinates energy, industry, compute, identity, payments, and duration warehousing. Functionality and collateral stability improve while user custody and practical exit contract. This state is not collapse; it is the strongest strategic competitor to the open stack.

Transitions are conditional. Quiet long-dated inflation compensation keeps State 1 available; a sustained rise associated with physical bottlenecks moves the system toward State 2. Failed repair plus coercive balance-sheet rules indicates State 3. Demonstrated closed-stack service delivery with shrinking user exit indicates State 4. None of these observations supplies a tactical forecast for bonds, currencies, or commodities.

A cross-cutting observable: household transmission. The four states are sovereign-level descriptions, and each has a household-level counterpart that can diverge from it — a state can be stable at the sovereign level while its households deplete, and the divergence is a transition signal in its own right (“The Vibecession Was Real”; Sources). The observables, all public series: the saving rate; delinquency by income cohort rather than in aggregate; first-time-buyer share and payment-to-income for *new* entrants against incumbents; underemployment and degree-required employment share; mean versus median net worth; and sentiment against spending (a persistent gap between dollar-weighted activity and person-weighted assessment is the divergence itself, not noise). The reason this belongs in the state model rather than only in §29’s adversary analysis: the states describe who receives sovereignty, and household depletion is how a state loses the political capacity to remain in its current state. A State 1 with exhausted household buffers transitions differently than a State 1 with intact ones — toward

whichever policy regime the depleted median will vote for, which is rarely the market-repair state.

27.3 Four Phases of Adoption

We sketch the curve as four overlapping phases, each pulled in by a different constituency. These are **metric-anchored, not calendar-anchored**—transitions happen when telemetry shows they have happened, not when a date arrives:

Phase I: Cypherpunk-led. Plumbing proves out publicly; loop works without permission.

Phase II: Demand-led. Budgeted workloads close the fee loop; capacity becomes a line item.

Phase III: Composability-led. Receipts become default integration; invisible infrastructure.

Phase IV: Policy-led. Receipts become recognized evidence; jurisdictional differentiation.

These phases overlap and vary by geography and sector. The pattern is robust: early idealists, then necessity-driven budgets, then invisible defaults, and finally institutional and policy recognition. **Calendar horizons are intentionally omitted**—they age badly and invite cheap dunking. What matters is whether the phase gate metrics (below) are met.

27.3.1 Phase Gates and Failure Gates

The following table makes the adoption curve **auditable**. Each phase has entry gates (observable thresholds) and failure gates (conditions that falsify the phase or regress to a prior phase).

Phase	Gate Metrics	Threshold
I → II	VERIFYPRICE exists for ≥ 3 canonical workloads with ≥ 6 months history VERIFYSETTLE for ≥ 1 corridor with refund_safe = 1.0 VerifyReach published for major regions Receipt volume non-trivial	p95 within SLO success $\geq 95\%$ succ ₁ > 0.7 ≥ 1000 receipts/day
II → III	Fee coverage $\geq 30\%$ of security budget ^a Budgeted workloads $\geq$ speculative in volume VERIFYSETTLE survives ≥ 1 policy shock Capex/OPEX lines reference verification	Sustained 6+ months Visible in receipt tags SLOs hold Public disclosures exist
III → IV	Collateral/reserve usage in multiple venues VERIFYPRICE/Settle stable through macro stress Duration-neutral holding cohort Fee coverage $\geq 50\%$	≥ 3 independent SLOs hold during crisis Measurable in on-chain data Sustained
IV	Legal codification of receipts Public-sector participation Geographical differentiation	≥ 2 jurisdictions Visible in dashboards Measurable in VerifyReach

Table 27.1: Phase entry gates for adoption curve.

Phase Entry Gates:

Calibration note on the 30% gate. The II→III fee-coverage threshold sits roughly an order of magnitude above anything achieved by the closest precedents (Filecoin, Render, Akash, Golem, Livepeer), whose native fee revenues have covered low single digits of issuance through their operating lives (Section 31.1.2). The number is not derived from those systems’ best outcomes; it is set at the level below which the security budget is issuance-funded and the asset is economically a service token. It should be read as a deliberately hard target with a known failure rate, and Red Line 5 is its enforcement mechanism. ^a*Coverage here is FeeCoverage as defined in Section 23.6.5 — retained fees over budget, burns excluded — not a fees-plus-burn sum; the historical label “fee+burn” carried over from earlier versions and is retired.*

Condition	Consequence
VERIFYPRICE p95 exceeds SLO for core workloads for ≥ 3 months with no credible remediation	Phase regression; SoV thesis weakens
Corridor refund_safe < 100% on admissible route (repeated, not isolated)	Phase regression; Private Money claim fails
Verification concentrates > 70% on single profile/jurisdiction for ≥ 3 months	Phase regression; decentralization claim fails
Receipt data becomes unverifiable (datasets unavailable, dashboards dark)	Telemetry capture; thesis unverifiable
Fee coverage collapses and workload mix becomes > 80% speculative	Phase II → I regression

Table 27.2: Failure gates for adoption phases.

Failure Gates (Phase Regression or Thesis Falsification):

Phase regression logic: Phases can regress. If Phase II metrics collapse (fee coverage falls, budgeted workloads disappear, corridors fail stress tests), the system slides back to Phase I. This is not failure of the technology—it’s failure of adoption. The telemetry makes it visible so stakeholders can respond.

27.3.2 Adoption Curve To Expect

The adoption curve for this new monetary substrate will not be a single “flip” moment. It will unfold in discernible phases, each with its own constituencies, failure modes, and telemetry.

The through-line is simple. Under Bretton Woods II, you ran money by buying reserves, hiring lawyers, and trusting gatekeepers. Under Verification, you run money by buying capacity: privacy capacity, proof capacity, verified FLOPs. **Budgets stop paying for promises and start paying for work that anyone can check.**

“Phase I–IV” are not just vibes. Each phase has rough, observable thresholds in the same metrics we have already defined (VERIFYPRICE, VerifyReach, VERIFYSETTLE, FERs, Work Credit utilization). The point of the curve is that you can tell, from dashboards and receipts,

which world you are in.

27.3.3 Phase I: Cypherpunk-Led (Prove the Plumbing in Public)

Phase I belongs to the people who already live in the future: open-source cryptographers, hardware hackers, privacy-chain communities, and early AI+ZK builders. Their job is to prove that the loop (Create/Compute → Prove → Settle → Verify) can be made real **without permission**.

In this phase, the main artifacts are reference implementations and receipts:

- **Open-admission prover markets** stand up on top of existing chains and zk networks. They publish live VERIFYPRICE dashboards for canonical workloads: PROOF_2²⁰, MATMUL_4096, basic PoL fingerprints.
- **MatMul-PoUW testnets** (Duplex-style patterns) test whether verification overhead can remain $r(W) = v/p \leq 0.3$ on commodity hardware and whether miners without favoured hardware profiles can still win blocks. This is a target under investigation, not a demonstrated result: the optimal-security claim remains conjectural (Table 20.3), and Phase I exists to produce the measurement rather than cite it.
- **PoL pilots** (Ambient-like patterns) demonstrate hybrid verified inference with honest-output rates that can be measured and contested, not marketed.
- **zk-PoW networks** (Nockchain-style patterns) act as public receipt ledgers (places where proofs from many domains can be anchored and timestamped).
- **Non-custodial BTC↔XMR/ZEC swaps with refund-safe UX** move from GitHub curiosities to tools that people actually use.
- **Federated e-cash mints** (Fedimint-pattern) extend Bitcoin’s privacy and settlement into community-scale banks: each federation acts as a Layer-5 node, converting Bitcoin into blinded e-cash with full privacy for internal transactions and Lightning for cross-federation settlement. The more federations exist, the harder it becomes to surveil or co-opt any single one.

(Note: “Duplex-style,” “Ambient-like,” and “Nockchain-style” are referenced as design patterns, not endorsements of any specific project or ticker.)

Nothing is “mainstream” yet. Most users are still speculators and hobbyists. But a few things become hard facts rather than hopes: you can buy proofs as a service from permissionless markets; you can pay for them over privacy rails without custody; and anyone with a laptop can verify the receipts.

Phase I telemetry and triggers: Phase I is real (not hypothetical) once:

- **VERIFYPRICE exists for a small set of canonical workloads.** At least a handful of public markets publish VERIFYPRICE(W) dashboards for $W \in \{\text{PROOF_2}^{20}, \text{MATMUL_4096}, \text{INFER_LM_7B}\}$ with months of history, and p95 times stay within stated SLOs under stress. (Large-model inference enters this set only when ZKML maturity permits; Section 20.10 gates it. Expect-

ing 70B-class verification on laptop hardware in Phase I would trip Red Line 1 at launch and thereby test nothing but the checklist.)

- **VERIFYSETTLE is measured for at least one serious corridor.** For some $BTC \leftrightarrow XMR$ or $BTC \leftrightarrow ZEC$ corridor C , $VERIFYSETTLE(C)$ is public and hits targets like $\text{success}(C) \geq 0.95$, $\text{refund_safe}(C) = 1.0$.
- **VerifyReach is non-degenerate.** At least one verifier network publishes $\text{VerifyReach}(N, R)$ for major regions with real multi-path reachability.
- **Receipt volume is non-trivial.** Receipt ledgers anchor a steady flow of PIDL receipts per day (≥ 1000) across multiple workloads.

Once those metrics are visible, the question is no longer “can this exist?” but “can it scale?”

27.3.4 Phase II: Allocator-Led (Proofs and Privacy Become Budget Lines)

Phase II begins when **allocators** (fund managers, corporate treasuries, exchanges, and large web platforms) start to treat privacy and proofs the way they once treated bandwidth: as recurring operating costs, not science projects.

By this point (**after Phase I gate metrics are met**, not by calendar date):

- Proof and compute networks publish open $VERIFYPRICE$ and reliability dashboards with historical data. You can see, month by month, how p95 verification times, costs, and failure rates behave under load.
- The PAL SDK and settlement adapters have been integrated into serious applications: analytics pipelines, custody stacks, compliance tooling, provenance layers for media.
- Privacy corridors (particularly $BTC \leftrightarrow XMR/ZEC$) have polished GUIs, reference libraries, and documented latencies. Refund failures are statistical outliers with post-mortems, not routine hazards.

Allocators do what they always do when facing repression and technological change: they **reclassify**. Instead of framing privacy and proof capacity as speculative tokens, they treat them as line items required to keep operating:

- A bank’s AI risk model must be run on verifiable compute with audit-friendly receipts, because regulators now ask for them.
- A media platform must attach cryptographic provenance to high-stakes content, because the liability of not doing so is too high.
- A trading venue must use private settlement rails to avoid leaking its entire order book and client graph.

They do not buy these capacities because they have converted to cypherpunk ideology, but because compliance and risk management now require math, not memos. Under Verification, “do nothing” is no longer the conservative option; it is reckless.

Phase II telemetry and triggers: Phase II is real once:

- **Fee coverage reaches a meaningful slice of security budget.** For at least one serious PoUW/proof network, fees tied to real workloads cover $\geq 30\%$ of miner/prover revenue over 12–24 months (FeeCoverage convention: retained fees, burns excluded).
- **Budgeted workloads dominate speculative ones in volume.** In receipt analytics, the *number* of proofs purchased by enterprises grows steadily, even if speculative volume remains higher in nominal terms.
- **VERIFYSETTLE stays inside SLOs through at least one policy shock.** For at least one high-volume corridor, VERIFYSETTLE remains within bounds across a visible policy or regulatory event without catastrophic failure.
- **Capex/OPEX lines reference verification capacity explicitly.** At least some institutions treat Work Credits as a line item (verification spend, privacy spend), not as a speculative asset bucket.

At that point, the triad is no longer “crypto” from the allocator’s perspective; it is infrastructure they **have** to pay for.

Two allocator paths, only one of which is monetary. “Allocator-led” is ambiguous, and the ambiguity is dangerous. Institutions can arrive in two very different ways, which look similar in a price chart and nothing alike in the telemetry.

Phase II-A — native allocator adoption. Institutions acquire the base asset; stake it or post it as collateral; pay native fees; use private settlement; purchase proofs or verified compute; participate in native liquidity; and hold self-custodied or directly redeemable claims. This is **monetary adoption**, and it moves every series on the Value Capture Board.

Phase II-B — wrapper-led financial exposure. Institutions purchase spot ETFs and treasury-company equity, use futures and options, and hold omnibus custodial claims, gaining price exposure without native use. This is **market adoption**. It may precede monetary adoption, seed it, or permanently substitute for it (Section 5.6); which of the three is an empirical question answered only by VerifyFlow (Section 24.8).

Revised Phase II gate. Assets under management, market capitalization, price level, and wrapper launch count **do not** qualify the system for Phase III. Passing requires positive native fee growth; material native burn or retirement; growing collateral lockup; expanding native settlement; rising receipt volume; healthy VERIFYPRICE, VerifyReach, and VERIFYSETTLE; a Wrapper–Native Growth Gap below a pre-declared threshold; and no uncontrolled custodial concentration.

False Phase Transition: Wrapper-Led False Positive

Market capitalization and institutional AUM rise rapidly while native use, fee coverage, burns, and collateral remain flat. The ecosystem appears to have advanced to institu-

tional adoption but remains, economically, in the speculative phase. This is the most flattering way for the thesis to fail, which is exactly why it needs a named gate rather than a footnote.

27.3.5 Phase III: Composability-Led (The Stack Disappears into Infrastructure)

Once Phase II gate metrics are met and sustained, composability becomes second nature. Developers no longer think in terms of “ZK project X” or “PoUW chain Y”—they think in terms of the loop.

Provenance, verified compute, and private settlement start to resemble TLS on the web:

- New applications default to emitting claims and receiving receipts via the SDK, simply because it is easier than hand-rolling trust.
- Major frameworks and toolchains ship with verification modules and privacy adapters bundled: verifying a receipt feels as ordinary as opening an HTTPS socket.
- Chains and rollups routinely outsource heavy computation to PoUW or proof factories and anchor receipts on shared ledgers.

From the outside, nothing dramatic happens. There is no “flipping.” What changes is the **default**:

- Sensitive data is processed either under enclaves with open silicon profiles or under ZK, and accompanied by attestations.
- Payments that cross borders or touch politically exposed persons quietly route over privacy rails, with receipts that satisfy auditors but not censors.
- AI systems whose outputs matter are either run with PoL-style verifiable inference or backed by helmets of audits and shadow runs, with failure rates visible rather than buried.

Networks that can demonstrate a clear, measurable link between work and value, paired with credible scarcity and decentralization telemetry, are the ones that establish **durable cash-flow accrual** to the asset. That is what this phase can show. Whether a monetary premium sits on top of it is a separate question, decided by the holder-side flow of [Section 11.10.8](#) rather than by the accrual, and not one that price alone answers in either direction ([Corollary 11.1](#)).

Phase III telemetry and triggers: Phase III is real once:

- **Non-trivial native collateral and reserve usage.** The base asset appears as collateral or reserve holdings in ≥ 3 independent native venues, while Work Credits remain separately reported service claims. The Native Monetary Buyer Map shows a visible self-custodied, loss-bearing cohort rather than inferring reserve status from lockup.
- **Macro sensitivity flips.** Price and flows react as much to changes in workload budgets (AI capex, compliance requirements, settlement volume) as to crypto-native news.

- **VERIFYPRICE and VERIFYSETTLE stability through macro stress.** During interest-rate shocks or liquidity crunches, the underlying utility, capacity, and safety metrics do not collapse.
- **Duration-neutral base-asset holding behavior.** Holder data shows a meaningful self-custodied cohort holding the base asset for agency and SoV properties, with holding periods measured in years; Work Credit inventory is reported separately as service demand.
- **Fee coverage $\geq 50\%$.** Sustained over multiple quarters.

When these show up in the dashboards and market structure, the base asset has advanced as a monetary candidate. The triad has demonstrated service relevance; Work Credits have not changed instrument class.

27.3.6 Phase IV: Policy-Led and Path-Dependent (Verification as Public Good)

Once Phase III gate metrics are sustained through macro stress, the curve becomes more path-dependent and political. If the stack delivers on its risk promises (decent decentralization, cheap verification, robust privacy corridors that are visibly abuse-resistant), states and institutions will begin to treat parts of it as **public goods** rather than threats.

Some jurisdictions will move first. They will:

- Codify cryptographic receipts as acceptable evidence in courts and regulatory filings.
- Mandate provenance proofs for certain classes of media or AI systems rather than ad hoc labelling.
- Recognize privacy rails with viewing-key regimes as compliant infrastructure rather than as dark pools.

Others will resist, preferring the comfort of chokepoints and legacy gatekeepers. That is where the anti-repression design matters. If privacy rails remain non-custodial and censorship-resistant, if proof and compute markets remain globally addressable, capital and talent can route around laggard jurisdictions the way data routed around telcos.

In the best case, the monetary role of the triad becomes self-reinforcing:

- Each wave of repression (negative real yields, capital controls, information crackdowns) pushes more savings and more workflows onto rails where verification is cheap and permission is irrelevant.
- Each wave of adoption increases the depth and liquidity of proof and compute markets, which in turn makes it cheaper and more obvious to use them for new domains.

If the triad earns a store-of-value premium in this scenario, it is not because everyone suddenly shares a philosophical vision, but because the world has learned, through trial and error, that capacities which are verifiable, censorship-resistant, and continually demanded are safer long-term anchors than promises that can be administered or revoked. This paragraph describes a best case, not a forecast, and the premium in it remains the unsized quantity of [Section 11.10.8](#).

Phase IV telemetry and triggers: Phase IV is real once:

- **Legal codification of receipts and corridors.** Multiple jurisdictions explicitly recognize PIDL-style receipts as valid evidence in regulation and courts; some codify acceptable VERIFYSETTLE or provenance requirements into statute.
- **Regulatory reliance on cryptographic proofs.** Policies and enforcement actions reference proof primitives (ZK proofs, PoL audits, receipt formats) rather than merely “records” or “reports.”
- **Public-sector participation.** Public institutions (development banks, treasuries, public broadcasters) use triad-aligned rails and publish their own receipt and corridor metrics.
- **Geographical differentiation.** A measurable share of volume and capital migrates toward jurisdictions and networks that respect lawful privacy and verifiable compute, as seen in regional splits in VerifyReach, settlement volume, and Work Credit usage.

This adoption curve is not guaranteed. Verification cost can creep; hardware or router oligopolies can re-centralize control; regulation can pinch on- and off-ramps harder than expected. But if we anchor the stack in verifiable machines, keep verification cheap and public, treat privacy as infrastructure rather than vice, and force our own claims through telemetry and receipts, this is the curve we can plausibly aim at.

27.3.7 The Negative Case: Useful Infrastructure Without Money

It is important to describe what it looks like if the triad remains useful infrastructure but never becomes money:

- Proof markets exist and grow, but fees are paid in stablecoins or fiat. The native asset has weak demand.
- Privacy corridors work, but most users route through custodial services that accept traditional payment.
- Verified compute is valued, but hyperscalers dominate and price in USD. Decentralized alternatives remain niche.
- The base token trades as a speculative asset correlated with crypto markets, not as a store of value with independent demand.
- Fee coverage never reaches 30%; the security budget depends on issuance subsidies indefinitely.

In this scenario, the *stack* succeeds (it delivers real privacy, proofs, and compute) but the *monetary thesis* fails (the native asset does not capture the economics). This is the “utility-token trap” outcome. It is not catastrophic—the infrastructure is still valuable—but it means the store-of-value claim was wrong.

The telemetry makes this outcome **visible**: fee coverage, native-asset fee share, bypass-channel indicators, and holder behavior all signal whether the monetary thesis is tracking or not.

This scenario is not hypothetical. It maps directly to one pole of a live debate among Bitcoin practitioners: the position that Bitcoin succeeds as an escape hatch and self-custody tool for a small elite, but does not disrupt the existing power structure or impose a free market. The thesis's value-capture conditions and telemetry regime exist precisely to make this distinction measurable rather than a matter of temperament.

27.4 Failure Gates

Phases can regress. If Phase II metrics collapse, the system slides back to Phase I. The failure-gate conditions and consequences are stated once, with the phase entry gates in [Section 27.3](#); they are not repeated here. The regression logic is worth restating because it is the part that gets forgotten: regression is not failure of the technology—it is failure of adoption, and the telemetry exists so that stakeholders can see the difference.

Chapter 28

Risk Analysis & Failure Modes

The risks to this emerging triad are not merely technical; they are structural, political, and economic.

28.1 Technical Risks

Every guarantee in this thesis reduces, at the bottom, to mathematics running on matter that somebody manufactured. That is where the first class of risk lives — not in the cryptography, which is either sound or it is not, but in the supply chain underneath it.

Proof system failures. A SNARK/STARK is broken or exploited. The blast radius is bounded before it happens: multi-ZK support, proofs tagged with system and parameter IDs, and migration paths mean a broken curve deprecates one proof family rather than the entire attestation layer.

Hardware capture. A major fab or TEE platform ships a backdoor. This is the Layer 0 argument wearing different clothes: profiles, lot sampling, and open hardware alternatives exist so that one compromised supplier cannot silently become the substrate of record for every receipt (Section 15.4).

Protocol bugs. Consensus bugs, bridge flaws, privacy leaks. Mitigation: reference implementations, staged rollouts, incident response.

The common thread is that none of these risks is existential on its own; each is existential only if it arrives *unobserved*. A broken proof system caught in staging is an engineering incident. The same break discovered two years after deployment, with every receipt since built on it, is a monetary event.

28.2 Economic Risks

The economic risks are quieter than the technical ones, and more dangerous for it. They do not announce themselves as failures; they arrive as slow drift that every individual reading can explain away.

Verification cost creep. If verifying proofs ceases to be much cheaper than producing them,

the entire asymmetry collapses — and it collapses gradually, one mildly disappointing benchmark at a time. Shows up as rising $r(W)$ in VERIFYPRICE; the trip wire exists because no single reading looks fatal.

Centralization. When specialized hardware or closed routers dominate, neutrality erodes. Shows up in concentrated facility IDs, rising entry latency.

Fee death-spirals. If demand falls, fees collapse, security budgets shrink. Mitigation: conservative base issuance, Work Credits encoding long-term demand.

Paying the guards in fresh scrip. The fee-coverage question deserves the concrete form, because the table of 30/50/80% targets is easy to read as a KPI and forget. A system that relies entirely on issuance to fund security is a warehouse paying its guards each morning in shares the warehouse prints that morning. The guards keep showing up — right up until the shares stop buying lunch. The fee-coverage trajectory is simply the schedule for switching them to wages.

28.3 Political Risks

Political risk is the one class this thesis cannot engineer around, only prepare for. The adversary here is not a bug or a market; it is an institution with jurisdiction, patience, and every incentive to frame the stack in the least flattering available light.

Bans & sanctions. Jurisdictions may declare privacy assets illegal, sanction contracts, or criminalize the use of certain clients. Mitigation: jurisdictional diversity (labs, foundations, hardware profiles, LPs); client modes that degrade gracefully (offline, mesh, sat-links); legal defense resources; clear separation between core protocol and specific front-ends so that UI bans do not equal protocol death.

Info-ops and framing. Media and governments can frame the stack as “criminal tech” or “national security threat.” Mitigation: lawful-privacy narratives grounded in receipts and compliance primitives; visible legitimate use cases (payroll, provenance, AI verification, critical-infrastructure audit trails); and an insistence on evidence (“show the receipts”) rather than slogans.

Regulatory chokepoints at the edge. App stores, banks, and ISPs can be pressured to block access even when the protocol is neutral. Mitigation: the Layer-1/2 work described earlier (alternate transports, side-loading, content-addressed and offline distribution) so that no single storefront, bank, or carrier can become a kill switch.

Political risk cannot be “engineered away.” It must be **distributed and prepared for**: diversified jurisdictions, many independent implementations, multiple access paths, and a culture that expects (and drills) for attempted bans and smear campaigns rather than treating them as unthinkable. The stack is built so that being declared unwelcome in one jurisdiction is

survivable; no layer assumes it will never be declared unwelcome anywhere.

28.4 Spec Drift, Vaporware, and Scoreboard Capture

Spec drift and vaporware can hollow out trust long before censorship does. Whitepapers without shipped code or inflated performance claims corrode the ecosystem’s credibility. In metrics, this shows up as:

- a widening gap between claimed and measured $\text{VERIFYPRICE}(W)$;
- a mismatch between Work Credit issuance and actual workload utilization;
- chains with high market caps and low receipt volume;
- clients that are nominally “live” but barely process real workloads.

Active liveness tracking (clients, explorers, throughput under real load, settlement safety metrics) should replace social metrics or TVL as the industry’s scoreboard. Projects that publicly document testnet-to-mainnet milestones and publish SLOs and incident reports exemplify the discipline needed.

No dashboards, no trust. If VERIFYPRICE , swap success and refund safety, decentralization telemetry, and corridor health are not public, treat claims as unpriced risk. A system that refuses to show its receipts is asking you to underwrite the very soft guarantees this thesis is designed to escape.

28.4.1 Scoreboard Capture in the Market Realization Plane

Scoreboard capture has a financial-product form as well as a protocol form, and it is the more seductive of the two. A product industry under competitive pressure increasingly asks “what ticker might catch a flow?” rather than “what does the investor actually need,” with capital chasing realized performance and sponsors launching ever more specialized wrappers. The result is a scoreboard on which a wrapper can be a triumph while its holders are destroyed — reverse splits keep the share price presentable, new inflows replenish assets, and fees keep accruing to the sponsor (Section 24.8.4).

Scoreboard Capture, Extended

Scoreboard capture occurs when product AUM, market capitalization, launch count, or trailing returns substitute for measurement of whether *capital survived*, *native services were used*, and *the monetary loop functioned*.

The countermeasures are the Capital Survival Ratio, the Wrapper–Native Growth Gap, and Native Use Share (Section 24.8). None of them can be gamed by launching another product.

28.5 Collateral Architecture Under Wrong-Way Risk

Collateral design chooses a distribution of failure; it does not eliminate one. At least three models must be reported and stress-tested against the same drawdown and DVC scenarios:

Pure-native collateral. Operators post only the base asset. This maximizes direct lockup and value capture but also maximizes wrong-way risk: price decline raises required units, liquidation or exit releases float, and falling capacity can weaken the service case that supported the asset.

Mixed collateral. Operators post a published mixture of native and external reserve collateral. Direct native capture is weaker, basis and custody risks enter through the external leg, but collateral coverage can remain more stable through a native drawdown.

Native plus insurance. Native stake supplies incentive alignment while a separately capitalized first-loss, insurance, or resolution pool absorbs defined tail losses. The pool must publish capital, exclusions, correlation, claims priority, and replenishment rules; an unfunded promise is not insurance.

The comparison must report collateral coverage against realized slashing exposure, liquidation sensitivity, operator survival, released float, DVC, and service-SLO recovery. Governance may choose among the models only after publishing those trade-offs. It may not describe the strongest value-capture design as the safest prudential design.

28.6 Red Lines: When the SoV Thesis Fails

The thesis is falsifiable. If any of these conditions persist without credible remediation, the store-of-value claim is **no longer defensible**. Severity labels are used consistently below and mean specific things: **kills** = the condition falsifies the thesis outright (the chain's hinge is severed); **weakens** = the condition removes one leg of support such that the thesis survives only in reduced form and must be restated. The label appears in each red line's closing paragraph and is not a tone choice.

Red Line 1: Verification Affordability Breaks

Condition: Physical VERIFYPRICE p95 exceeds SLO bounds for core workloads for ≥ 3 consecutive months, with no credible remediation path.

Why it kills the thesis: "Anyone can verify" is the hinge. If verification becomes expensive, proofs become platform claims, not public facts.

The touchstone is the image to keep. Gold earned its monetary career on a dark stone against which a nugget could be streaked in seconds — expensive to smelt, cheap to assay, and impossible to fake convincingly. The entire premium of this thesis rests on the assay staying cheap. The moment verification costs real money or real expertise, receipts revert to what

platform labels already are: assertions with better typography. Three months of sustained breach is not a performance regression; it is the anchor reverting to the broadcast era.

Red Line 2: Refund Safety Breach

Condition: Corridor refund_safe < 100% on any admissible route, with repeated incidents and no automatic delist + remediation.

Why it kills the thesis: Non-custodial settlement is the “Private Money” foundation.

One lost deposit is an incident. A corridor that eats funds and remains listed is a verdict: the system has told every user what it actually is. Refund safety is binary by design — there is no partial custody, no 97% non-custodial — because trust does not survive being mostly honored. The automatic delist matters more than the breach itself; a network that keeps a broken corridor on its board for volume’s sake has chosen the scoreboard over the user.

Red Line 3: Verification Monoculture

Condition: > 70% of verifications on a single hardware profile, TEE vendor, or jurisdiction for ≥ 3 consecutive months.

Why it kills the thesis: Monoculture means “trust the dominant vendor.”

Three-fifths of the world’s proving running on vendor-attested chips from one supplier is not an efficiency; it is a single point of political failure wearing a supply-chain dashboard. The concentration can arrive without any malicious act at all — one vendor simply being better for two quarters is enough — which is precisely why the line is drawn on the reading, not on intent. The moment verification monocultures, the adversary no longer needs to break the cryptography; it needs only to lean on the vendor.

Red Line 4: Telemetry Capture

Condition: Receipt datasets become unavailable, unverifiable, or controlled by a single party.

Why it kills the thesis: If telemetry can be captured, the entire observability regime is theater.

A constitution nobody can check is a press release. Every instrument in Part V exists on the assumption that an independent party can re-derive the published number from raw public artifacts and get the same answer; the moment the raw data itself becomes a proprietary good, the dashboards continue to render but stop meaning anything. This is the quietest of the kill conditions — no price moves, no users are harmed, nothing visibly breaks — and it is the precondition for every other red line’s measurement. Capture the scoreboard and every other game on it is rigged.

Red Line 5: Fee Coverage Collapse

Condition: Fee+burn coverage falls below 10% of security budget and workload mix becomes $> 80\%$ speculative for ≥ 12 months.

Why it kills the thesis: The SoV story requires structural demand, not pure narrative.

This is the guards-in-fresh-scrip condition (see [Section 28.2](#)) promoted from metaphor to trip wire. A network can coast on issuance for years, and during those years every headline number looks healthy: hash rate, market cap, even transaction counts. What it cannot do is fake the composition of its demand. When four out of five units of workload exist to farm a token rather than to be used, the fee floor is not being tested — it is being performed, and the security budget is a subsidy awaiting its first serious drawdown.

Red Line 6: Value Capture Failure

Condition: Triad usage grows (proof volume, settlement volume, compute demand) but native-asset fee volume, burns, collateral lockups, and fee coverage do not grow with it for ≥ 12 months.

Why it kills the thesis: The system may be useful infrastructure but not a store-of-value asset. Users are consuming triad capacity through bypass channels.

The paper-gold precedent is the cautionary case, and it is a century old: claims on gold outgrew holding of gold, exchanges multiplied paper while vaults stood still, and the price of the paper discovered it could move without the metal underneath it. The same substitution is available here, and the telemetry exists to watch it arrive. Usage growing while native fee volume stagnates is not a paradox — it is the market routing around the asset while continuing to use the network, and it is the single clearest signal that the thesis's monetary link has quietly severed while the infrastructure succeeded.

Red Line 7: Legal Incompatibility

Condition: Lawful users in major jurisdictions cannot use privacy rails without unacceptable compliance uncertainty for ≥ 12 months, and no lawful-privacy patterns (viewing keys, scoped disclosure) are adopted.

Why it weakens the thesis: If regulated actors cannot participate, the “Private Money” leg loses its institutional constituency and anonymity sets shrink.

Privacy that only outlaws can use is a shrinking set, not a monetary property. The thesis's lawful-privacy design exists precisely to avoid the corner where compliance uncertainty meets default privacy and the only remaining users are the ones who were never complying. Anonymity is a network good: each legitimate departure thins the crowd that makes blending plausible, and the set collapses gradually — which is why this line weakens rather than kills,

and why its telemetry (corridor compliance receipts, viewing-key usage) reads on adoption, not on law.

Red Line 8: Governance Capture

Condition: Governance can alter issuance schedules, fee routing, or telemetry rules without hard constitutional constraints, timelocks, or supermajority requirements — *or* the narrow emergency path (Section 23.5.3) becomes routine rather than exceptional, with its subtractive-only scope eroded to permit any value redirection.

Why it kills the thesis: If insiders can redirect value away from holders or inflate supply at will, the asset is a platform token, not a store of value.

The constitution here is whatever governance cannot silently violate without getting caught (Section 23.1 states the principle). A body that can only shrink the protocol cannot steal from it — that is the entire logic of the subtractive-only emergency path, and the reason its erosion is a kill condition rather than a governance hygiene note. The failure arrives as precedent, not as heist: one emergency invoked for convenience, one scope quietly widened, and the constraint that made the asset safe to hold has been converted into a knob.

Red Line 9: Wrapper Dominance Becomes Monetary Substitution

Condition: Wrapper Dominance Ratio (Section 11.13) rises for ≥ 2 consecutive quarters while native fee share, private settlement volume, and collateral lockups stagnate or decline. The WDR is a *pair* — its stock component (custodial-exposure share of total exposure) and its flow component (wrapper-activity ratio) — and “rises” means *both components* rise over the window; if one rises while the other falls, the condition is not met and the split is published. In its full form: custodial and synthetic exposure (CCR, SER) becomes the dominant form of ownership for a sustained period while protocol-native fees, collateral, private settlement, proof demand, and verified-compute usage remain stagnant or declining, with a persistently positive Wrapper–Native Growth Gap in its unit-denominated form (Section 24.8) — the USD form is published alongside but is not the trip wire, because it can move on revaluation alone.

Why it weakens the thesis: Price rises while the monetary thesis dies. Government-approved wrappers becoming the dominant institutional form is co-option without censorship. The asset may remain a successful financial product, but the claim that it functions as a native store of value *for Privacy, Proofs, and Compute* is no longer supported by anything except its chart.

Red Line 10: Physical Infrastructure Opacity

Condition: Physical VerifyPrice (Section 20.8) exceeds threshold, or Facility Capacity Receipt data is unavailable or unverifiable for a material share of active capacity, for ≥ 2

consecutive quarters.

Why it kills the thesis: Layer 0 claims become trust-me claims; Work Credits lose infrastructure credibility.

Red Line 11: AI Enclosure

Condition: Top-3 hyperscalers, closed TEEs, or a single jurisdiction exceed 50% of VERIFYPRICE-tracked verified compute capacity (the 40% chip-family/TEE cap of [Section 15.4](#) being the early-warning tier of the same concentration measure), i.e., the Enclosure Risk Flag ([Section 5.6](#)) triggers and persists.

Why it kills the thesis: The verified-compute service market described by the “AI Money” lens becomes a cloud IOU; the frontier is enclosed rather than homesteaded, so it cannot support the associated base-asset monetary candidate.

Red Line 12: Agency Failure

Condition: Forced Disclosure Incidence ([Section 25.1.2](#)) becomes systemic, or participation/agency use cases ([Section 3.6](#)) fail to materialize while institutional usage grows, for ≥ 12 months.

Why it kills the thesis: Lawful privacy collapses socially, or the protocol serves institutions but not users—violating the ninth SoV requirement ([Item 9](#)). Emergency governance becoming routine is a related failure: neutrality decays into foundation fiat.

Red Line 13: Energy Sovereignty Failure

Condition: Network capacity-weighted sovereign optionality $\mathcal{O}_s^{\text{net}}$ ([Section 15.8](#)) falls below threshold for ≥ 2 consecutive quarters, **or** $> 50\%$ of VERIFYPRICE-tracked capacity sits in jurisdictions operating an active curtailment or power-rationing regime against verification workloads.

Why it kills the thesis: Verification affordability becomes a sovereign policy variable rather than a market outcome. Red Line 1 becomes *externally triggerable*: an adversary or host state can break the hinge by adjusting tariffs, interconnection queues, or load priority, without touching the cryptography, the governance, or the market.

Red Line 13 and its neighbours. Three red lines now touch the physical substrate and they are not redundant:

- **Red Line 10** is *opacity*: we cannot see the substrate.
- **Red Line 13** is *fragility*: we can see it clearly and it will not survive pressure.
- **Red Line 11** is *enclosure*: we can see it, it is robust, and someone else owns it.

A network can breach any one without the others. Perfect FCR disclosure of a single-interconnect fleet in a rationing jurisdiction breaches 13 while passing 10. Do not merge them.

The dependency this makes explicit. Section 20.7 describes Physical VerifyPrice SLOs as constitutional and exogenous to token price. That is right, and Red Line 1 correctly treats their sustained breach as fatal. But exogenous to token price is not exogenous to everything: the SLOs remain endogenous to the physical and jurisdictional conditions under which power and reference hardware are obtainable (Section 5.5).

Red Line 13 exists so that this dependency is monitored upstream rather than discovered downstream. Without it, the sequence

energy policy $\rightarrow$ proving cost and siting $\rightarrow$ Physical VERIFYPRICE $\rightarrow$ Red Line 1

runs to completion with no instrument reading anywhere along it until the hinge itself fails. A falsification framework whose primary condition can be tripped by an unmonitored external variable is not yet falsifiable in the way it claims.

Red Line 14: The Capturable Wedge Closes

Measured quantities. For canonical workload W at tier T in period t :

- $P_{W,T,t}^{\text{prot}}$ — all-in USD price to the buyer of executing W at tier T through the protocol, inclusive of the protocol fee. Protocol-native: PIDL receipts already carry workload ID, tier, timestamps, and hardware profile (Section 20.3).
- $P_{W,t}^{\text{byp}}$ — lowest publicly quoted USD price to execute the same computational content *without* protocol-grade verification, on the same hardware class. Sourced from a **pre-committed reference panel**: a frozen list of venues (e.g., three hyperscalers and three GPU marketplaces), scraped daily, with the workload-to-instance mapping published as a machine-readable spec. Panel changes require publication and 90 days' notice. Where a venue publishes committed-use, reserved, or spot pricing, the *lowest generally available* of those is the reading, not the on-demand list price.
- $s = (P^{\text{prot}} - P^{\text{byp}}) / P^{\text{byp}}$ — **normalized bypass spread**. Normalization is not cosmetic: the absolute spread shrinks mechanically with compute deflation, so only the ratio is economically meaningful.
- γ — verification cost share: proving plus redundancy per unit, expressed as a fraction of P^{byp} . Derived from FER and VERIFYPRICE telemetry, and subject to the same third-party custody and reproducibility requirement as the reference panel (below), because γ enters ω negatively and the fee recipient has a direct interest in understating it.
- $\omega = s - \gamma$ — the **capturable wedge**: what is left of the price the protocol commands after paying for the verification that justifies it.

- τ — **realized take rate**: *gross* native fees (USD) / (settled protocol turnover, USD). Gross, not net-of-burn, because τ measures what buyers pay — the same convention the FeeCoverage definition pins for the fee family (Section 23.6.5); burns are supply arithmetic, read as reduced net issuance, and adding them would count destruction as payment. Computable by anyone from chain data.

Condition A — the wedge closes. Capacity-weighted ω across the canonical starter set (Section 20.10) falls below the realized take rate, $\omega < \tau$, for ≥ 2 consecutive quarters, with no credible remediation path.

Condition B — the base deflates faster than volume compensates. Native fee turnover fails to grow over a trailing four-quarter window while protocol physical throughput (verified units delivered) grows over the same window. “Fails to grow” means the trailing-four-quarter sum is non-increasing ($\leq$ its prior value) for two consecutive quarters. Deflation adjustment is by the *workload unit-price index* — the capacity-weighted average USD price of the canonical workload basket — not by the asset price: the condition asks whether buyers are transacting less even as physical delivery grows, so the deflator strips compute-price deflation, not asset-price revaluation. Both the raw and deflated series are published. **Persistence:** the red line reads on the condition holding for two consecutive quarters, not one; a single flat quarter is a flag, not a trip.

The connective, stated once: either condition breaching trips the red line. A breach of A, or a breach of B, or both, is a breach of Red Line 14. There is no requirement that they breach together.

Why A kills the thesis: The protocol is charging more than the differential value it supplies. It is living on switching costs, subsidy, or inertia rather than on anything a buyer would pay for, and the fee is not a wedge on genuine differential value but a tax on captive volume. That fee is retractable by competition, so nothing durable accrues.

Why B kills the thesis: The wedge may be intact and the burn still cannot scale. An *ad valorem* fee on a deflating unit price is a shrinking real toll (Section 11.10.2); if throughput grows while real fee-plus-burn turnover does not, the empirical bet the thesis is making on volume outrunning deflation is losing, in public.

Both clauses must be monitored, because each closes a gaming route the other leaves open. A published alone is gameable by cutting the fee toward zero, which trivially restores $\omega > \tau$ while capturing nothing; B published alone is gameable by inflating volume at a vanishing wedge. Monitoring both means a protocol that is neither differentially valuable nor growing its real take has nowhere to stand. This is a requirement on the *test*, not on the trigger: dropping either clause from the instrument panel is a breach of the falsification protocol, while breaching either clause in the readings is a breach of the red line.

Why Condition A is hard to game. It is **self-normalizing**. It compares two independently measured quantities — what the market will bear net of verification cost, and what the protocol actually takes — rather than testing either against a threshold somebody had to guess. There is no parameter to lobby for. A protocol cannot pass by lowering its ambitions, because lowering the take rate lowers τ and lowering the price lowers s , and the comparison survives both.

Third-party checkability. The thesis's own principle applies with full force here: the party being scored must not set the scoring function (Section 15.8, Red Line 4). The reference panel, the workload-to-instance mapping, and the verification cost share γ must therefore all be maintained by someone other than the fee recipient, and an independent party must be able to re-derive ω from the published inputs and obtain the same number. γ deserves explicit mention because it is the easiest term to shade: it enters $\omega = s - \gamma$ negatively, so understating the cost of verification inflates the wedge and makes this red line harder to trip, and the fee recipient is the party with both the incentive and the telemetry. Its inputs — proving overhead and redundancy per unit, drawn from FER and VERIFYPRICE — must be published in the same reproducible form as the panel, with the derivation runnable by an outside party against the raw receipts. A wedge computed by the entity collecting the wedge is a marketing figure, and that applies term by term.

Honest coverage caveat. P^{byp} is well defined only where a like-for-like unverified execution exists: matrix multiplication, inference, proof generation. It is *not* well defined for atomic settlement or media provenance, where there is no unverified version of the same product — the comparator there must be the all-in fee of a custodial or regulated rail delivering the same economic function. That is a rougher measurement with a weaker claim to like-for-like, and it must be reported separately rather than blended into a single index. An index that averages a clean comparison with a rough one inherits the rough one's error and hides it.

Effective prices, not list prices — and the direction of the residual bias. On-demand list pricing is not what hyperscaler and marketplace capacity actually transacts at. Committed-use discounts, reserved instances, and spot markets routinely clear well below list, so a panel scraping list prices measures a bypass channel nobody uses. That is why P^{byp} is specified above as the lowest *generally available* price across the published pricing modes rather than the on-demand quote, and why the panel spec must record which mode each reading came from.

Even so, a residual bias survives: the largest buyers transact at negotiated enterprise rates that are published nowhere, so the observable price remains an upper bound on the true alternative. The direction this pushes the red line is worth deriving rather than assuming, because it is not the direction one expects. Since $s = (P^{\text{prot}} - P^{\text{byp}})/P^{\text{byp}}$ falls as P^{byp} rises, an overstated P^{byp} understates both s and γ , and understates s by more, because P^{prot} exceeds the

per-unit verification cost. Measured ω is therefore too small and Condition A trips *earlier* than the truth warrants. The reading is accordingly published as an **upper bound on breach**: a measured $\omega < \tau$ may overstate the case against the thesis and must be re-derived against negotiated pricing before the 90-day clock is treated as running, whereas a measured $\omega > \tau$ is a pass earned against a conservative comparator.

That is the benign direction, which is precisely why the term that runs the other way needs naming. Understating γ inflates ω and makes the red line harder to trip, and γ is the one input sourced from the fee recipient's own telemetry. The custody requirement above exists for that reason and is the load-bearing safeguard here, not the panel methodology.

Red Line 14 and its neighbours. Three distinctions worth stating precisely, because each is easy to collapse:

- **The price series is not VERIFYPRICE.** VERIFYPRICE is a *cost* SLO — the real-resource cost of checking a claim on reference hardware, deliberately exogenous to token price (Section 20.7) — and Red Line 1 reads on it. Red Line 14 does not: the quantity it reads on is a *price* series for delivered service, set by markets, which VERIFYPRICE is not and was never meant to be. The two red lines therefore cannot be tripped by the same event. The qualification is that γ , one input to the wedge, does draw on VERIFYPRICE and FER telemetry for the cost of verification, which is the right source for a cost term; that is why γ carries the independent-reproducibility requirement above. Shared telemetry in one term, not a shared test.
- **This is not Red Line 6.** Red Line 6 asks whether fees follow usage — a direction. Red Line 14 asks whether there is a wedge worth charging at all — a magnitude (Section 11.10.4). A protocol can pass 6 while failing 14: native fees can track usage perfectly while the fee itself exceeds the differential value supplied.
- **This is not Red Line 5.** Fee coverage measures fees against the security budget. Red Line 14 measures the fee against what the market will bear. Coverage can be adequate on a base that is about to be competed away.

Red Line 15: Native Collateral–Capacity Spiral

Measured quantities. Publish native-asset peak-to-trough drawdown, collateral coverage against realized slashing exposure, operator exit and liquidation rates, stress-scenario DVC, collateral released into liquid float, and recovery of core service SLOs.

Condition. A pre-declared material native-asset drawdown persists for at least 30 days; collateral coverage falls below realized slashing exposure; operator exit or forced liquidation exceeds its constitutional threshold; stress-adjusted DVC declines; released native collateral materially increases liquid float; and core service SLOs fail to recover within the 90-day remediation period. **The clauses are conjunctive by design** — each alone has an innocent explanation (a drawdown without coverage failure is a bear market;

operator churn without DVC decline is turnover), and it is the joint deterioration that names the spiral. The thresholds must be set before the episode, reported for pure-native, mixed, and native-plus-insurance designs, and may not be relaxed while the clock is running.

Why it kills the thesis. The asset is being offered as insurance while its underwriting capacity disappears in the insured event. The loop

drawdown → more units required → operator exit → DVC decline → collateral release → liquid float inc

turns native value capture into a service-capacity contraction. A temporary drawdown is not the breach; the sustained joint deterioration and failed SLO recovery are.

Pattern: The red line is not “bad thing happens once.” It is “**sustained breach + no recovery.**” Isolated incidents with rapid remediation are expected in any complex system. Persistent degradation without response is thesis failure.

Condition	Response
Red line breached	Incident declared; governance must publish remediation plan within 14 days
Remediation fails after 90 days	Asset reclassified from “SoV candidate” to “speculative/experimental” in protocol communications
Multiple red lines breached simultaneously	Crisis mode; independent review commissioned; results published

Table 28.1: Red line response protocol. The clocks run on distinct triggers and do not nest: the 14-day publication clock starts at incident declaration; the 90-day remediation clock starts when the remediation plan is published (or at day 14 if none is); the two overlap only by construction of that sequencing, and each is reported with its own start date on the incident page. Neither clock resets on parameter change — the no-relaxation rule of Red Line 15 generalizes: thresholds and clocks are frozen for the duration of the episode they measure.

What happens at a red line: These are not punishments—they are **truth in advertising**. An asset that claims SoV properties must demonstrate them. If it can’t, it should stop claiming.

28.7 Market Realization Warnings

Red lines are protocol failures. The Market Realization Plane generates a second, weaker class of signal that must not be confused with them. A wrapper unwind can destroy price without damaging protocol function at all; treating that as thesis falsification would be as sloppy as treating a rally as confirmation. The distinction:

- **Monetary red lines** (above) mean the protocol's monetary claim has failed: verification breaks, settlement safety breaks, native value capture fails, issuance becomes discretionary, censorship routes dominate, or native use can be bypassed.
- **Market realization warnings** (below) mean *price has become an unreliable signal*. They do not necessarily kill the protocol; they suspend the right to cite price as evidence about it.

Warnings trigger when:

1. A statistically significant break appears in holder flow elasticity.
2. Net mechanical gain κ changes sign.
3. The Mechanical Pressure Ratio exceeds pre-declared market-depth bounds.
4. Wrapper creations dominate native spot demand.
5. Dealer swap capacity appears to bind.
6. Exposure migrates from swaps toward options.
7. Recursive leveraged wrappers appear.
8. Top wrappers or dealers exceed concentration thresholds.
9. The Wrapper–Native Growth Gap stays elevated.
10. Price rises while native fees and receipt volume fall.
11. Price falls while native monetary health improves.
12. Cross-jurisdiction hedges produce halt or settlement asymmetries.
13. Triad capacity is consumed at scale while the asset prices as a service and the monetary bid accrues elsewhere — **Service-Good Realization**, below.

The correct response to a warning is not intervention. It is **epistemic**: publish the warning, state which flow term is dominating, and stop using price as evidence in either direction until it clears. Governance must specifically not respond to warnings by inflating issuance, subsidizing price, or buying back supply — those are the reflexes of an entity managing a stock price, not operating a monetary constitution.

One warning is not merely epistemic. The twelfth item above is the last of the flow-mechanical warnings. The thirteenth is a different animal and is set out in full, because it is the observable form of the thesis losing the argument of [Section 31.1.1](#) and because a warning that carries falsifying weight should not be left as a line in a list.

Warning 13: Service-Good Realization

Condition. All three clauses hold together, read across at least two distinct *regime-pressure episodes* — intervals in which the Stage A index R_t ([Section 11.14.1](#)) is elevated under criteria published before the reading rather than chosen after it:

1. **Capacity is being consumed at scale.** Receipt volume, verified units delivered, and private settlement volume grow through the episode; the capturable wedge stays above the realized take rate, so Red Line 14 does not fire; native fee-plus-burn

turnover tracks usage, so Red Line 6 does not fire. *The protocol is working.*

2. **The insurance signature is absent.** Through the episode the asset exhibits nothing that distinguishes it from a claim on a service business: holding duration does not extend, the non-custodial share does not rise, the duration-neutral holding cohort of [Section 27.3](#) does not grow, and demand moves with buyers' capacity budgets rather than with their exposure to the pressure. Read on the Value Capture and Agency Preservation boards ([Section 24.9](#)), not on a chart.
3. **The monetary bid accrues to Bitcoin.** Over the same episodes Bitcoin absorbs the flow the thesis predicts for a regime-contingent hedge, while the triad asset covaries with compute and IT spending rather than with regime pressure.

What it means. Not that the protocol failed — clause 1 says it did not. It means the regime-contingent convenience yield of [Section 11.10.8](#) was not there, and that yield is the entire remaining basis for the monetary claim once [Section 11.10.7](#) has conceded that fees, burns, and collateral produce a cash-flow claim and a floor. The prescribed response is therefore *reclassification rather than retirement*: the asset is a verified-capacity service asset with a competitively priced fee stream and a collateral floor. That is a real thing to be. It is not a store of value, and the protocol's communications must stop saying otherwise on the same 90-day clock the red-line protocol uses.

Why this remains a warning rather than another red line. Clause 3 reads on price, and reads it against another asset. Every condition in [Section 28.6](#) is protocol-observable, and [Corollary 11.1](#) forbids treating price as evidence about monetary adoption in either direction. A price-comparative red line would break both commitments. Red Line 15 instead covers the protocol-observable collateral–capacity spiral. What distinguishes this warning from its twelve neighbours is that clauses 1 and 2 carry the falsifying weight by themselves: capacity consumed at scale with no insurance signature, across repeated episodes, is the monetary claim failing whether or not anyone looks at Bitcoin. Clause 3 says where the premium went. It is corroboration, not the test.

Chapter 29

Implementation Sketches for Builders

Everything above is the blueprint. This section is about what to actually ship first.

29.1 Minimum Viable Stack (MVS) Sequencing

Builders will ask: “What is the minimal sequence?” The answer matters because dependencies are real.

Priority	Component	What It Enables	Dependency
1	PIDL + verifier harness + anchoring	Telemetry cannot be faked	Foundation
2	One corridor kernel + VERIFYSETTLE	Non-custodial settlement	PIDL
3	One proof factory + VERIFYPRICE	Proof commodity market	PIDL
4	PAL SDK + PRK SDK	Developer surfaces	Corridor + factory
5	Neutral router + fairness tests	Open participation	SDKs + dashboards
6	PoUW pilots	Useful-work security	All above stable

Table 29.1: Minimum Viable Stack sequencing.

Why this order:

1. **Receipts first:** Without verifiable receipts, all claims are unpriced. Ship PIDL, reference verifiers, and dataset anchoring before anything else. This is the foundation that makes “no dashboards, no trust” real.
2. **One corridor before many:** A single refund-safe, non-custodial corridor with public VERIFYSETTLE is worth more than ten “coming soon” announcements. Ship $BTC \leftrightarrow ZEC$ or $BTC \leftrightarrow XMR$ with 100% refund safety before expanding.
3. **One proof factory before PoUW:** Prove that canonical workloads can be produced, verified, and priced before adding consensus complexity. PROOF_2²⁰ and MATMUL_4096 are good starting points.
4. **SDKs after primitives:** PAL and PRK are developer surfaces, not primitives. Ship them once the underlying corridors and proof factories are stable.
5. **Neutral routing after basic markets:** Fairness tests and admission metrics matter only once there’s enough activity to measure. Don’t over-engineer routing before you have provers

to route.

6. **PoUW last:** Proof-of-useful-work consensus is the most complex piece. It requires stable proof factories, receipts, and measurement infrastructure. Don't attempt it until Priorities 1–5 are solid.

Hardware profile clarity: “Commodity GPUs with well-understood drivers” is an “open-ish hardware profile.” To be explicit: this is **Profile Class: Partially Closed, High Familiarity**. GPUs are not verifiable machines in the same sense as open RTL—they are “known black boxes” whose behavior is well-characterized but not inspectable. Label profiles clearly so users understand the trust assumptions.

29.2 Layer-0 Verifiable Machines

Start with small, fully open “security evaluation” chips: key-storage/signing cores, simple TRNG/PUF arrays. For each tape-out, commit to a fixed sampling plan.

Minimal Layer-0/4 pattern:

- Pick one open-ish hardware profile (e.g., commodity GPUs).
- Define a hardware profile HID with chip family, driver versions, entropy tests, power metering.
- Stand up a proof factory cluster with reference prover binaries.
- Publish basic telemetry and VERIFYPRICE samples.

29.3 Privacy Rails: Making Settlement Safe and Boring

Standardize adaptor-signature swap flows. Take $\text{BTC} \leftrightarrow \text{XMR}$ and $\text{BTC} \leftrightarrow \text{ZEC}$ as canonical corridors. Write down message formats, time-lock conventions, refund procedures.

Ship wallet UX with recovery, not bravado. Clear progress indicators, explicit timeouts, visible “abort and refund” buttons.

Document settlement latencies and success rates. Publish historical p50/p95 time-to-finality, success rates, and failure breakdowns.

29.4 Proof Factories: Receipts as a Service

The next implementation target is the Prove stage: turning raw work into portable receipts that any chain, app, or institution can consume.

Unify proof outputs under a receipt schema. Start by defining a compact, chain-agnostic receipt format (PIDL): the claim hash, workload ID, circuit or model hash, proof hash or tran-

script commitments, SLA tier, start/end timestamps, resource usage, and provider signatures. Whether the underlying artifact is a MatMul transcript, a logits fingerprint, or a SNARK, all of them fit into this envelope.

Expose a single programming surface. For developers, the SDK should offer one set of verbs: `declare_claim`, `request_proof`, `await_receipt`, `verify_receipt`. Underneath, it can target multiple proof systems and networks. The developer doesn't wire to "Platform X"; they attach their claim and policy to the SDK and let it find anyone who can satisfy the SLA.

Integrate multiple zkVMs and proof systems from day one. Monoculture is the enemy of both security and economics. The factory should be multi-ZK and multi-backend by construction: support at least one pairing-based SNARK, one STARK or FRI-based system, and one zkVM at launch, with a clear path to adding more.

Make telemetry a first-class product. Every proof request and receipt should emit metrics into the VERIFYPRICE Observatory: p50/p95 verifier times and costs for each workload, failure modes, queue depths per backend, variance under adversarial mixes.

29.5 Compute Consensus Pilots

With privacy rails and proof factories in place, you can begin to anchor consensus itself in useful work. Think in terms of reference pilots, not one chain to rule them all.

MatMul PoUW pilot (Duplex-style): Launch a devnet where the block-making puzzle is a MatMul instance drawn from a canonical distribution. Use the workload registry/WF-ABI to describe tasks; derive instances from header randomness; implement low-rank noise so that miners can return both a correct product and a succinct proof. Wire block validity to the presence of a valid MatMul receipt whose VERIFYPRICE parameters are below published thresholds. Instrument everything: $r(W)$ ratios, verifier times on consumer hardware, time-to-first-proof for new miners, and centralization metrics.

Verified-inference pilot (Ambient-style PoL): In parallel, build a network where the work function is "serve model inferences with proofs of honesty." Commit to models and datasets; design logits-fingerprint schemes that are simple, deterministic, and uniform; add randomized audits and peer-prediction-based slashing. Here full ZKML may be too expensive, so hybrid verification is acceptable, but only if it is exposed honestly: publish honest-output rates, audit coverage, and failure patterns as SLOs, not marketing.

zk-Proof-of-Work / receipt ledger pilot (Nockchain-style): Finally, either integrate with or prototype a chain whose primary job is to produce and timestamp zk proofs themselves. This network can serve as a common receipt ledger: proofs from many domains are anchored here

with minimal metadata and availability guarantees.

These pilots need not be over-promised as final destinations. Their job in the early years is to act as laboratories: to establish that verification asymmetry holds in practice at scale, that decentralization can be preserved under useful-work mining, and that receipts remain cheap to check.

29.6 Developer Kit: Make “Import Proofs” the Default

The last implementation sketch is about developer experience. If every application needs a team of cryptographers to participate, nothing scales.

Claims & proving library: A library that plugs into common languages and frameworks and exposes a simple API: mark this function as “must be proven,” mark this data stream as “must be provenance-tracked,” specify acceptable backends and SLAs, and let PAL handle task generation and proof requests.

Settlement adapter: A component that connects to existing wallets and key-management setups, and that exposes “pay-for-proof” and “pay-for-compute” as simple intents: no manual swap logic, no bespoke bridges, just a predictable interface plus receipts, backed underneath by PRK and the atomic-payout kernel.

Verification module: Modules that can be compiled into smart contracts, browser bundles, and command-line tools, so that end-users and auditors can check receipts without standing up new infrastructure. `verify(receipt)` should be as embedded in the tooling as `log()` is today.

The work of the next few years is, in large part, to make these three pieces boring. When “import proofs” feels like “import TLS,” and when “settle privately” feels like “call the payments API,” the cypherpunk monetary stack stops being an argument and starts being the default way serious systems are built.

Chapter 30

The Closed Sovereign Stack

Everything argued so far has been argued in the abstract. The thesis asserts that when soft guarantees weaken, value migrates toward hardened, verifiable, vertically coherent capacity. It then spends six parts describing how to build such a stack.

A reasonable reader will ask whether anyone is actually doing this, and whether it works.

The answer is yes, at civilizational scale, by a state, with the politics inverted. That case deserves a chapter of its own—not as a prediction, an endorsement, or a geopolitical forecast, but because it is the closest thing this thesis has to an existence proof and to a counterexample simultaneously.

Epistemic Status of This Chapter

The material here is of three kinds and they are not interchangeable.

Primary sources, used for what they state: published policy speeches and central-bank staff research. Where this chapter quotes an official, the quotation is verbatim and the claim is about stated doctrine, not about outcomes.

Analyst commentary, principally Doomberg’s *Fire Horse* presentation and macro work by Luke Gromen. Neither is peer-reviewed. Figures attributed to them are *their* accounting rather than independent measurement, and are used to characterize a strategy rather than to establish a quantity.

Scenarios, which are quarantined in [Appendix I](#) together with the claims from the same sources that did not survive checking.

No claim in this chapter is load-bearing for the thesis. If the figures are wrong in detail, the argument here weakens as illustration and the rest of the document is unaffected. Nothing in [Section 28.6](#) depends on it.

30.1 A State That Already Assumes Soft Guarantees Have Failed

Chapter 5 builds an adversary model from the premise that institutional promises are becoming conditional. That premise is usually defended by argument. It can also be defended by observation: some actors are already behaving as though it were true, and paying enormous

costs to do so.

China's energy and industrial strategy is best read as a sustained bet against the reliability of external guarantees. The recurring pattern is not accumulation of any single resource but maximization of the number of independent pathways from a controlled input to a required output:

- Domestic coal retained at scale, not because it is clean or cheap in the long run, but because it can be mined inside the system and therefore cannot be interdicted at sea.
- Refining capacity built beyond immediate domestic need, with feedstock flexibility across crude grades, LPG, and ethane.
- Coal-to-liquids, coal-to-gas, and coal-to-chemicals facilities that are difficult to justify on conventional return-on-capital grounds.
- Nuclear technology licensed from multiple mutually hostile suppliers, plus indigenous designs, so that no single provider controls the program.
- Caution about additional pipeline dependence even on friendly suppliers.
- Statutory minimum coal stockpiles—which purchase not fuel but *time*, the interval during which the system can adapt.
- Electrification of transport, converting an imported-oil vulnerability into demand for electricity and batteries produced domestically.

The organizing principle is conversion optionality: many paths, deliberately redundant, accepting poor measured returns in exchange for freedom of action. Several of these investments are straightforwardly value-destroying under normal market conditions. That is the point. Their value is state-contingent, concentrated in exactly the scenarios where ordinary market relationships stop clearing.

A synthetic-fuel plant that loses money nineteen years in twenty and prevents a transportation crisis in the twentieth is a bad investment and an excellent insurance policy. Discounted cash flow prices the first sentence and not the second.

Resilience buys time, and time is a weapon. The stockpile item above is easy to read as a defensive measure, and reading it that way understates it. A statutory minimum inventory purchases an interval. What the holder does with that interval is a separate question, and the answer is not necessarily defensive.

A system that can defer, substitute, ration, or absorb a disruption acquires something more useful than protection: it acquires influence over *when the disruption becomes binding*. It can decline to bid during a scramble and replenish afterward, at lower prices. It can wait out an adversary's electoral cycle, funding round, or refinancing calendar. It can allow a shock to run knowing that the counterparty with the shorter inventory and the more leverage-sensitive balance sheet will have to move first. Redundancy converts into tempo, and tempo converts

into bargaining power:

resilience → time → strategic initiative

Strategic Tempo

The conversion of physical resilience—stockpiles, spare capacity, fuel-switching and conversion flexibility—into control over *when* a disruption becomes binding, and thence into bargaining power. Tempo is the return on redundancy that conventional cost accounting does not price.

This is a strictly stronger claim than the option-value argument of [Section 15.10](#), and it compounds with it. Option value says the redundant facility pays off in the bad state. Tempo says its existence changes which bad states occur, and when—because a counterparty that knows you can wait negotiates differently than one who knows you cannot. For Layer 0 this means the sovereign optionality index of [Section 15.8](#) is measuring something whose value is understated by its own construction: the index counts pathways, and pathways buy time, and time is worth more than the pathways cost.

This is the same structure as the option value argued in [Section 15.10](#): facility value contains commercial, contingency, deterrence, and bargaining components, and conventional analysis captures only the first. A country that can survive an embargo has more leverage *before* the embargo, whether or not the capacity is ever used.

Why this matters for Premise 1. The thesis argues in [Section 1.2](#) that soft guarantees are weakening. A state spending decades and enormous sums to become less dependent on external promises is expressing the same judgment through capital allocation rather than through argument. Revealed preference at sovereign scale is a stronger form of evidence than commentary, and it is available to the thesis without requiring any normative agreement with the actor.

30.2 Trust Minimization Applied to Matter

The cypherpunk instruction is familiar:

Don't trust; verify.

The strategic instruction described above is structurally identical:

Don't trust; possess, duplicate, stockpile, or retain an alternative conversion pathway.

Both are responses to the same problem—a promise is weaker than a possession—and both accept significant efficiency losses to reduce dependence on counterparty good behavior:

- A contractual assurance that oil will arrive is weaker than an inventory of oil.
- An assurance that a vendor will keep supplying technology is weaker than domestic production.
- An assurance that an ally stays friendly is weaker than diversified supply.
- An assurance that trade lanes remain open is weaker than domestic conversion capacity.

Every line above is a Layer 0 argument ([Chapter 15](#)) written in molecules instead of silicon. The recognition is useful in both directions: it tells us the thesis’s physical premises are not eccentric, and it tells us that trust minimization is politically neutral machinery that serves whoever builds it.

30.3 The Mirror: The Same Doctrine, Arriving in the West

The preceding sections treat the closed stack as a foreign object—coherent, instructive, and elsewhere. That framing has an expiry date, and it has probably passed.

On 23 June 2026, the U.S. Secretary of the Treasury told the Economic Club of New York that “economic security begins with national capacity,” identified semiconductors, artificial intelligence, quantum computing, advanced manufacturing, shipbuilding, critical minerals, and pharmaceuticals as the industries that would define the next century, and argued that supply chains must be able to withstand coercion and crisis rather than merely minimize cost. The speech quotes Hamilton directly on enlarging the sphere of domestic commerce. The Hamiltonian framing is not an outside interpretation imposed on the policy; it is the policy’s account of itself.

Read against [Section 30.1](#), this is recognizably the same doctrine: capacity valued above cost, resilience valued above efficiency, strategic continuity valued above return on capital.

What the speech does not say, and the discipline of noting it. It does not advocate gold settlement. It does not propose capital controls. Its stated posture is “open to the world while anchored at home,” and it treats digital assets, stablecoins, tokenization, and payment standards as instruments for extending dollar financial leadership rather than for retreating from it. Commentary that reads this material as a program of monetary retreat is reading in something that is not present, and [Appendix I.4](#) lists the specific claims of that kind that this thesis declines to use. The accurate description is selective protection plus industrial subsidy plus investment screening plus dollar-denominated digital rails—which is a profound regime change without being the one the more excitable commentary describes.

The asymmetry. What makes the Western version harder than the original is not doctrine but the balance sheet it must be financed from.

The closed stack of [Section 30.1](#) was built while suppressing household consumption, directing credit administratively, and tolerating poor measured returns for decades. The Western attempt begins from the opposite configuration and must preserve, simultaneously: high asset valuations, high corporate margins, entitlement commitments, low consumer prices, an open capital account, a reserve currency, and low inflation. These cannot all be maximized at once, and [Section 3.2](#) explains why the constraint binds harder than it looks—the asset valuations are not merely a political preference but a fiscal input.

The West is attempting to move from financial sovereignty back toward physical sovereignty, financed through a balance sheet optimized for the regime it is leaving.

There is a second, independent reason the Western attempt is hard, and it is Green's rather than Gromen's. China's directed banking system is a **duration warehouse**: it can be compelled to hold the interval between pouring concrete and producing power even when expected private returns are poor. That creates misallocation and hidden losses. It also enables projects whose strategic option value exceeds their commercial value. The United States has relied on foreign insurers, liability-driven pensions, retirement defaults, market-value indexes, and leveraged relative-value intermediaries. [Section 3.1.2](#) argues that this buyer architecture is not a substitute for the old liability-matched bid. The physical reconstruction therefore arrives at the same moment the financial system is becoming less willing to own maturity. The closed stack solves time by command. The open financial system is currently solving it by shortening the public book into bills—which transfers duration onto the rollover calendar rather than extinguishing it ([Section 3.2](#)).

The reshoring trilemma. Reindustrialization can be fast, cheap, or excellent; a system may choose two. Strategic competition demands fast. Bond yields and inflation demand cheap. The security rationale is void unless the result is excellent. After decades of deindustrialization the binding constraints are not primarily financial—skilled trades, engineering depth, machine tools, grid interconnection, transformer and switchgear supply, permitting throughput, supply-chain density, and industrial project management are all lead-time constrained, and capital cannot compress lead times past a point. The realistic expectation is therefore reshoring that happens, and takes longer and costs more than its advocates project.

Reindustrialization is a restructuring of claims. This is the part with monetary consequences, and it is the reason the chapter belongs in this document rather than in a policy annex.

Converting a financialized economy into a productive one means turning financial claims into factories, generation, transmission, mines, machine tools, trained labor, and fabrication ca-

capacity. If an economy has issued more long-duration claims on future output than its physical system can honor *while also* funding that conversion, the transition must impair some claims. The impairment can arrive as inflation, negative real rates, taxation, margin compression, capital controls, or explicit restructuring, and in practice it arrives as a combination.

The industrial transition is also a liability-side restructuring. Someone's claim on future output is not going to be honored at present value, and which someone is a political question with no automatic answer.

The tempting simplification is generational—bondholding retirees versus wage-earning young. It does not survive contact with the distribution. Inflation harms renters, cash-poor workers, households without bargaining power, and anyone trying to acquire assets before wages adjust; older households hold equities, real estate, and operating businesses, not only bonds. The defensible statement is narrower:

Reindustrialization redistributes away from the beneficiaries of asset-price inflation, cheap imports, suppressed labor shares, and long-duration financial claims—toward some combination of labor, domestic producers, strategic industries, and the state. Which combination is decided by policy design, not by the transition itself.

That last clause is the whole of the thesis's interest in the question. An industrial policy without broad ownership and wage gains does not dissolve a financial oligarchy; it substitutes a subsidized industrial one, and the participant is no better off for the change. This is [Section 3.6](#) restated at national scale, and it is the reason [Section 5.2](#) treats protective industrial policy as an adversary class rather than as a neutral backdrop.

30.4 Open Weights Are Not Open Sovereignty

The doctrine of [Section 30.1](#) has extended to intelligence, and the extension is easy for a reader of this thesis to misread as good news.

Openly released model weights are now within reach of the proprietary frontier. The Kimi K3 release, a 2.8-trillion-parameter mixture-of-experts model published with full weights, is presented by its authors as an open frontier system; the same paper reports that it trails the strongest proprietary models overall while leading the other systems in its evaluation suite. That is the claim this thesis makes and no more: open-weight capability is close enough to the frontier to be economically consequential.

The consequence is a familiar strategic pattern applied to a new good. In solar, batteries, electric vehicles, steel, and telecommunications equipment, the objective was never to capture the highest margin. It was to make adequate capability abundant at a price beneath the level

competitors' capital structures require. Software weights are unusually well suited to that strategy, because once trained they can be distributed at approximately zero marginal cost. A sector priced for durable scarcity rents does not need to be defeated technically; it needs only to have the scarcity assumption questioned credibly. Whether this is designed as economic strategy is not observable and is not claimed—[Appendix I.4](#) sets out why intent claims here are unusable. The effect does not depend on the motive.

The category error. None of this makes open-weight models an ally of the open stack, and the temptation to treat them as one should be resisted firmly.

Open Technology vs. Open Sovereignty

Open-source or open-weight software is sovereign only when its users can independently obtain the power, hardware, data, communications, privacy, and settlement required to operate it. Openness at one layer is fully compatible with concentration at every layer beneath it.

A model can be open at the weights layer and still depend on concentrated semiconductor supply, state-supported training runs no independent party could finance, centralized cloud for inference at useful scale, datasets that cannot be inspected or reproduced, national telecommunications, and alignment and provenance decisions embedded during training that no downstream user can audit. Published weights answer one question—may I run this?—and leave every question in [Chapter 15](#) untouched.

open model $\neq$ open sovereign stack

Open sovereignty additionally requires inspectable or verifiable models, portable execution, hardware diversity, privacy, local or distributed compute, independent energy and communications, user-controlled identity and settlement, and the practical ability to fork or exit. That list is the stack this document specifies, and the gap between it and an open-weights release is the entire subject of Parts III through V.

A gift at the weights layer is not a transfer of sovereignty. It may be the most efficient way yet devised to make a dependent population feel independent.

The practical instruction follows the same rule the rest of this chapter uses: judge a stack by what its users can do without permission, not by what its license file says.

30.5 The Inversion: Who Is Sovereign

The convergence is structural. The politics are opposite, and the difference is not rhetorical—it appears in specific mechanisms.

Dimension	American digital-dollar stack	Industrial-collateral stack	Open bearer stack (this thesis)	Mechanism here
Primary sovereign	The state, via regulated intermediaries	The state	The participant	Item 9
External trust	Reduced by protected domestic capacity	Minimized by state capacity	Minimized by verification	Chapter 15
Internal trust	Concentrated in licensed institutions	Concentrated at the center	Replaced by public checking	VERIFYPRICE
Privacy	Visible to regulators and law enforcement	Visible to authorized institutions	Private by default, selective disclosure	Viewing keys, Layer 5
Identity	Institutional and account-based	Administrative and persistent	Contextual and proof-based	Layer 3, PIDL
Settlement	Tokenized, dollar-denominated, permissioned at the edge	Permissioned, policy-responsive	Non-custodial, bearer-like	Layer 5, PRK
Compute	Strategic national champions	Nationally integrated	Portable, independently verifiable	Layer 4
Redundancy	Allied and reshored supply	Controlled from the center	Produced across independent operators	$\mathcal{O}_s$, Section 15.8
Duration warehouse	Bills, retirement defaults, leveraged intermediaries	Directed banks; state mandate to hold the interval	Not supplied by Layers 0–6; labeled credit beside duration-neutral money (Section 3.3 , Section 31.5.1)	Diagnosis, not a layer
Chief danger	Enclosure by rescue	Administrative enclosure	Fragmentation and weak value capture	Section 5.2 , Lemma 11.1

Table 30.1: Three sovereign stacks: convergent structure, different locus of sovereignty. The first two columns differ in method and agree that the state holds the keys.

Two readings of this table matter more than the rest.

The first is the last row. The open stack’s characteristic failure is not oppression but irrelevance: fragmentation, coordination failure, and value that accrues to wrappers rather than to the base asset. That failure mode is the subject of [Lemma 11.1](#) and of Red Lines 6 and 9. Every column has a way of losing, and the thesis is not entitled to compare its best case against another column’s worst.

The second is that the table has three columns rather than two, which is a change from how this comparison is usually drawn—including in earlier versions of this document. The convenient framing sets a closed foreign stack against an open Western one. The evidence in [Section 30.3](#) does not support it. Two stacks are currently under construction in which the state holds the keys; they differ in method, legitimacy, and the degree of legal recourse available to a participant, and those differences are real and worth a great deal. They do not differ on the

question this thesis asks, which is whether a participant may exit without permission.

30.6 What This Validates

Three claims in this document become harder to dismiss as utopian.

Hardened stacks are achievable. A frequent objection to Layer 0 is that verifiable machines, energy provenance, and jurisdictional dispersion are too expensive and too slow to be real. The closed stack demonstrates that vertically coherent physical infrastructure can in fact be built deliberately, at scale, over decades, against market signals. The question is who pays and who benefits, not whether it is possible.

Resilience is priced by serious actors. [Section 15.10](#) argues that redundancy carries option value invisible to normal cost accounting. That argument is not a rationalization invented to defend this stack's overheads. It is the operating assumption of states making the largest infrastructure allocations in the world.

The physical layer is monetary. The thesis's insistence that Layer 0 belongs in a monetary argument, rather than in an appendix about datacenters, matches how sovereigns actually reason about energy: as the substrate of the capacity to honor claims.

30.7 What This Warns

The warning is sharper than the validation, and it is the reason this chapter sits before the objections rather than in an appendix.

Abundant power and hardened infrastructure do not produce liberty. They produce capacity. Capacity can equip a population or condition it, and the same generating plant serves both.

A stack that is energy-abundant, industrially self-sufficient, computationally sovereign, and administratively integrated has everything required to make participation conditional: electrified transport, domestic payments, integrated identity, state-directed credit, national cloud, and pervasive telemetry. Nothing in the engineering resists this. The technology is identical; the difference is who holds the keys and who may exit.

This is precisely the participation line of [Section 3.6](#) and the enclosure risk of [Section 5.6](#), arriving not as speculation about a possible Western drift but as a functioning system. Red Lines 11 and 12 exist because the same capacities this thesis wants to build can be assembled

into the opposite arrangement, and the difference is not visible from a capability audit.

The unresolved question of the coming period is not whether systems become more integrated, computational, and hardened. They will. It is whether the people inside them are owners or tenants.

The thesis's normative purpose therefore cannot be “build a hardened stack.” That target is already being hit. Its purpose must be to build one whose hardness protects the participant rather than the administrator—which is a claim about key custody, exit rights, disclosure defaults, and value capture, not about throughput.

30.8 What We Do Not Claim

Discipline here matters, because this material is unusually easy to overread.

Energy consumption is not economic value. Large energy throughput demonstrates industrial scale. It does not establish productivity, capital efficiency, household welfare, or the quality of investment. A system can consume enormous energy while destroying capital, and redundancy financed by suppressed consumption and bad debt transfers costs rather than eliminating them.

Resilience is not automatically superior to efficiency. Excess capacity can become misallocation, local-government liability, environmental damage, and chronically weak returns. The question is never whether resilience is good but whether the insurance was correctly priced—which is exactly why [Section 15.8](#) insists on a measured index rather than a slogan.

The closed stack has not escaped its own constraints. A very large imported-oil dependence remains. Refining flexibility, stockpiles, electrification, and synthetic fuels buy time and bargaining power; they do not create domestic petroleum. The accurate description is increasingly *energy-resilient*, not *energy-sovereign*.

This is not a prediction. We are not forecasting which stack prevails, nor asserting that the closed model is ascendant. The claim is narrower and structural: both models are responses to the same collapse of soft guarantees, and they differ in who receives sovereignty.

This is not an endorsement. The strategy is analytically instructive and politically inverted relative to everything in [Chapter 4](#). Describing a system's coherence is not approving of it.

30.9 What This Changes in the Stack

A chapter that only reframed would not belong in Part VI. Four concrete consequences follow.

1. **The competitor is not fiat, and there is more than one of it.** The thesis is usually read as arguing against soft-guarantee fiat and custodial intermediation. The more serious competitor is a hardened, competent, permissioned stack that delivers verification, settlement, and compute with excellent uptime—and conditions access. [Section 30.3](#) indicates two such stacks are being built rather than one, by systems that regard each other as adversaries and that converge on the same answer to the question of who holds the keys. Both win on convenience and capability, not on principle. The thesis must therefore compete on exit rights and value capture rather than on capability alone, and it should expect no jurisdiction to be a natural home.
2. **Dispersion is a monetary property, not an operational preference.** If the alternative to a neutral stack is a competent closed one, then jurisdictional and grid dispersion stop being engineering hygiene and become the substance of the neutrality claim. This is what [Section 15.8](#) measures and what Red Line 13 protects.
3. **Capability audits are insufficient.** Two stacks can post identical VERIFYPRICE, uptime, and throughput while differing entirely in whether users can leave. Telemetry must therefore include agency and exit metrics ([Section 3.6](#), Red Line 12), not only performance. A scoreboard that measures only capability cannot distinguish the three stacks compared in [Table 30.1](#)—which is the most important thing it could tell us.
4. **Physical sovereignty requires a financial time horizon.** Energy determines what can be built; duration finance determines whether it remains funded until completion. The closed stack warehouses that interval by command. The American digital-dollar stack currently warehouses it with retirement defaults, market-value indexes, and leveraged intermediaries, or postpones it by issuing bills. The open stack specified here warehouses *claims* that are duration-neutral; it does not, and must not, pretend that Layers 0–6 are a pension. The completeness problem—who finances the reactor—is an objection, not a layer, and is answered in [Section 31.5.1](#).

The missing layer: household transmission. The comparison table above stacks sovereign capabilities against sovereign capabilities, and the stack this thesis defends can lose a comparison the table cannot express. Sovereign resilience and household resilience are different quantities, $R_{\text{state}} \neq R_{\text{household}}$: a state’s capacity to sustain energy, industry, defense, and administration can be intact — can even be improving — while the median household’s capacity to sustain shelter, savings, family formation, and shock absorption collapses, because the state’s continuity was partly financed by consuming the household’s buffers (“The Vibecession Was Real”; Sources). This is not a redistribution argument; it is a stability argument. A

household running a negative operating surplus does not stop consuming — it consumes out of stock:

$$O_i = \underbrace{Y_i + I_i + T_i}_{\text{income}} - \underbrace{E_i + DS_i}_{\text{essentials + debt service}}, \quad O_i < 0 \Rightarrow \Delta R_i < 0$$

where R_i is the household's resilience stock: liquid savings, unused borrowing capacity, accessible asset equity, reliable social insurance. Aggregate consumption stays smooth the entire way down, which is precisely why the aggregates are a lagging indicator of the exhaustion gradient of [Section 3.6](#) — dollar-weighted measures do not see person-weighted depletion until it surfaces as delinquency, and by then the buffers are gone. The relevance to this chapter's comparison is direct: the closed stack's strongest claim is that it *works*, and “works” measured only at the sovereign layer can mean the state's stability was funded by directing its households down that gradient. **Distributed resilience — R_{state} intact and $\text{Median}(R_i)$ not collapsing — is therefore a requirement the open stack must be able to evidence, not a social preference appended after the engineering:** a state that survives by transferring fragility downward eventually receives the fragility back through politics, defaults, labor quality, and the demand for control that §3.6 already established as the repression flywheel. The participation-line telemetry is where that evidence lives.

The joint failure mode: services from the closed stack, price from the wrappers. The two adversary chapters of this Part — the closed sovereign stack and the market-realization plane — are usually read separately, and the most likely way this thesis dies is their combination, which no single dashboard currently names. Run the causal chain: the closed stack matures into State 4 service delivery ([Section 27.2](#)) and supplies provenance receipts, compliant privacy predicates, and subsidized sovereign compute at quality and price the open stack cannot match for the median commercial workload — which is exactly the bypass channel of Red Line 6 and Condition A of Red Line 14, arriving through a public-sector door rather than a hyperscaler one. Simultaneously, wrappers and custodial products remain the only legally holdable form of the base asset in major jurisdictions, so whatever exposure demand survives is realized entirely through the market machine ([Section 11.14](#)). The end state satisfies every element of the thesis's own definition of failure while reading as success on most instruments: the stack's dashboards stay green because the open stack keeps functioning; the price holds or rises because wrapper demand is price demand, not protocol demand; and the base asset becomes a **reference price with no native loop** — quoted, held, and hedged, but never used to transact, prove, or exit.

Three readings of the joint scenario, in order of severity:

- **The red lines already contain it, separately.** RL6 and RL14 Condition A fire on closed-stack service substitution; RL9 fires on wrapper-led realization; RL12 fires on agency use cases failing to materialize while institutional usage grows. What none of them captures is the

joint condition — all three firing together is the coherent failure, and each firing alone has an innocent explanation (a hyperscaler price war; a bull market; slow enterprise adoption).

- **The diagnostic signature is a triple divergence**, readable before any individual red line trips: Wrapper–Native Growth Gap persistently positive (exposure outgrowing use), convenience-yield telemetry flat or falling across rising regime pressure (Section 24.8.6 — no one pays extra to part with units because the bearer service is not the one being demanded), and Homestead Ratio falling while closed-stack service metrics improve. Any one of these is noise; the triple divergence is the joint scenario arriving.
- **The thesis should not survive it quietly**. If the joint scenario materializes, the honest reclassification is the one Section 28.6 prescribes for any breach — the asset is a financial product and the stack is infrastructure, and the monetary claim is retired. The purpose of naming the mode here is that the triple divergence is observable *years* before the individual thresholds bind, and a thesis that publishes fifteen red lines should not be killable only by the one failure mode it failed to name.

The household amplifier: why the joint failure arrives non-linearly. The joint scenario above is a demand-side story, and its most likely trigger is an asset repricing that lands on a household sector with no buffers left to absorb it. Before the shock, resilient aggregates are the sum of two individually fragile mechanisms: upper-cohort discretionary spending out of asset gains, and lower-cohort necessity spending out of depletion. An equity or AI-collateral repricing removes the first term; the second cannot expand to replace it, because depleted households have no R_i left to convert into consumption — so aggregate spending does not decline gently, it breaks (Section 30.9). The same structure runs through the sovereign layer: upper-cohort spending carries tax receipts, so the shock propagates into deficits and Treasury stress, and the policy response faces the trilemma Section 27.2 is really about — tighten for inflation, ease for collateral, or build for supply, each worsening another constraint. For this thesis the implication is a measurement one: the base asset’s recession-hedging claim is regime-contingent on household buffers, not merely on asset correlations, and the convenience-yield telemetry (Section 24.8.6) is the instrument that would register whether holders actually pay for the holder-side service in exactly that state — the one state where the thesis says it matters.

Chapter 31

Objections & Responses

Any thesis that ends with “this becomes money” deserves hard pushback.

The objections below are grouped into five blocks. The **monetary** block attacks the central claim and is the one to read if you only read one. The **infrastructure** block disputes whether the stack can be supplied. The **political** block disputes whether power permits it, and whether we would remain distinguishable from what we oppose. The **market-structure** block disputes whether financialization quietly substitutes for monetary adoption. The **completeness** block asks whether duration-neutral money, even if it works, still leaves civilization unable to finance long-lived assets. Each objection is stated in its strongest form we could construct before it is answered, and several are conceded in part.

31.1 Monetary Objections

These are the objections that matter most, because they attack the claim the thesis actually makes: that a useful capacity can become a store of value. If any of the five below succeeds, the stack may still be worth building, but it is infrastructure rather than money.

31.1.1 “Why not just Bitcoin?”

Objection: Bitcoin already satisfies the store-of-value requirements this thesis specifies, and has satisfied them on a fifteen-year record under sustained adversarial pressure: credible scarcity on a schedule nobody can change, verification cheap enough to run on a laptop, censorship resistance demonstrated rather than modelled, bearer custody with no issuer, and no duration for policy to pin negative. It achieves this through a work function whose defining property is that **nobody outside the system buys the work**. That is not an inefficiency awaiting correction; it is the source of the objectivity. A puzzle with no external buyer has no customer to lobby, no supplier to subsidize, no regulator with jurisdiction over its demand curve, and no way for anyone to alter what it costs except by spending more of the same resource. Proof-of-useful-work deliberately introduces an external buyer and therefore weakens precisely that property. So everything in this document is either something Bitcoin already does, or a change that makes the money worse.

Response:

- **The monetary half of this objection is correct, and it is conceded without qualification.** Bitcoin's work function is monetarily superior to proof-of-useful-work, and superior for exactly the reason given. Unforgeable costliness with no external buyer is the strongest available form of the property, because the absence of a buyer is what makes the cost non-negotiable. Introduce a buyer of the work—an enterprise purchasing verified inference, a regulator requiring provenance receipts—and you have introduced a party who can be subsidized, taxed, mandated, prohibited, or captured, and whose demand can therefore be moved by something other than the resource cost of the work. **PoUW trades monetary objectivity for capacity relevance.** That is a trade, and this document is not entitled to call it an upgrade. Any account that describes SHA-256 as waste or heat is wrong, and wrong substantively rather than stylistically: it misidentifies the absent buyer as an inefficiency when the absent buyer is the product.
- **The thesis's claim is narrower than the objection assumes—and narrower than parts of this document have implied.** It is not that a better money is available. It is that Bitcoin does not, by itself, supply portable attestation of computation and provenance, or verified compute. It was never designed to, and there is no defect in its not doing so. Those are separate goods with separate buyers and separate demand curves: a bank's model-risk function buying proof capacity is not in the market for a monetary asset, and would not be satisfied by one. Whether an asset that supplies those goods accrues monetary premium, or merely trades as a priced service, is an open empirical question. The thesis has already conceded most of the ground here: [Section 11.10.7](#) states that fees, burns, and collateral establish a cash-flow claim and a balance-sheet floor rather than moneyness, and [Section 11.10.8](#) relocates the monetary claim to a convenience yield it declines to size. The honest position is that the question is open and the telemetry exists to answer it, not that the answer is known and favourable.
- **The slots Bitcoin occupies are not the slot this asset can compete for.** [Section 11.10.8](#) works this out and it is not restated here. In summary: gold holds “works when the network does not,” Bitcoin holds “works when the state does not,” and neither requires a grid at the moment of use. A triad asset requires power, reachable networks, and obtainable reference hardware in order to deliver anything at all, which confines it to a narrower band—adversarial conditions severe enough to need proofs and privacy, but not severe enough to take the grid. Part I argues that band is the likely one. Arguing it is likely is not the same as sizing it, and the thesis does not size it.
- **In this architecture the thesis is a consumer of Bitcoin, not a competitor to it.** Every privacy corridor specified at Layer 5 is a BTC↔ZEC or BTC↔XMR corridor ([Section 21.3](#)); VERIFYSETTLE is defined over those corridors ([Section 21.5](#)); the Class A trustless row of the bridge classification is an atomic swap with Bitcoin on one leg ([Section 21.4](#)); federated

e-cash mints extend Bitcoin's settlement rather than displacing it (Section 27.3); and receipt ledgers anchor to a neutral chain, of which Bitcoin is the obvious candidate. Bitcoin is also named as a bridge asset that survives the transition rather than a precursor to be discarded (Section 27.1.1). A document that wanted Bitcoin displaced would have to explain why it built its settlement layer on top of it. Damage to Bitcoin is damage to this stack.

- **What losing to Bitcoin would look like, specified so it can be checked.** The thesis loses this argument if triad capacity is consumed at scale—receipts, verified units, private settlement all growing, the capturable wedge intact, native fees tracking usage—while the asset shows no behaviour distinguishing it from a claim on a service business, and the monetary bid accrues to Bitcoin instead, across repeated episodes of the regime pressure the hedge is supposed to be for. That pattern is named and its measurement specified as **Service-Good Realization** (Section 28.7). It remains a market-realization warning because its price-comparative clause cannot be a red line without breaking Corollary 11.1. Red Line 15 instead covers the protocol-observable native collateral–capacity spiral.

The summary is uncomfortable and is better stated by us than by a critic: **Bitcoin is the better money; the triad supplies different goods, and whether those goods carry monetary premium is unsettled.** That is a weaker claim than a succession narrative and a considerably more defensible one, and it is what Section 11.10.8 already implies. Readers who came here expecting the thesis to argue that Bitcoin is superseded should note that no such argument is made anywhere in this document.

31.1.2 “This is just another chain / coin.”

Objection: We've seen this movie: every new system claims to be hard money with a fancy narrative. This is just buzzwords on top of a token.

Response:

- The core claim is not “this token pumps,” but “Privacy, Proofs, and Compute are verifiable necessities that can be priced and audited.”
- The focal metrics are **VERIFYPRICE/Reach/Settle and triad usage**, not just TVL or number of wallets.
- Work Credits are minted against **measurable work**: proofs, private settlements, verified FLOPs, with energy and hardware profiles tied in.

If, in practice, the system behaves like “just another chain” (no canonical workloads, no PoUW, no privacy corridors, no telemetry), then the objection is correct. The point of the architecture is to make that failure **obvious**, not to hide it.

The objection has an empirical record, and the thesis should cite it against itself. “Utility does not imply money” is not a conceptual puzzle; it is a decade of measured negative results. Filecoin, Render, Akash, Golem, and Livepeer each bootstrapped real infrastructure, real us-

age, and real fee-bearing demand — storage, GPU rendering, compute, and video transcoding respectively — and in none of them did native fees approach a meaningful share of the security budget, and in none did the asset acquire monetary behavior beyond beta to the sector. Fee coverage on the closest precedents has remained in the low single digits of issuance for years. This is the strongest available form of the objection, and it deserves the strongest form of the answer:

- **The precedents did not satisfy the conditions, and the thesis does not dispute the outcome.** Each of those systems fails or partially fails the [Value Capture Lemma](#) conditions as this thesis specifies them — fees payable in alternatives, minimal burns, weak collateral lockup, and (most decisively) service demand largely bypassable through centralized equivalents at similar quality. Their observed failure of premium accrual is what the lemma predicts, not a counterexample to it.
- **But the precedent cuts the other way too.** The II→III phase gate ([Section 27.3](#)) demands fee coverage of $\geq 30\%$ of the security budget sustained for six months — roughly an order of magnitude above anything the closest precedents achieved. The honest reading is not that the gate is modest but that it is deliberately set where no comparable system has reached, because below it the security budget is issuance and the asset is, economically, exactly what the objection says: a token on top of infrastructure. The gate should be read as the thesis betting against the precedent set, and Red Line 5 exists to settle that bet in public rather than around it.
- **What would distinguish this attempt.** Not narrative, and not telemetry theater: fee-bearing demand that cannot route around the asset (Condition 5, and see [Section 11.12.1](#) for the incumbent channel), verified at the workload tier ([Section 20.10](#)). Until the boards show it, the default expectation informed by the precedents is that this is another service token, and the burden sits on the design, not on the reader’s imagination.

31.1.3 “The triad does not belong in one token.”

Objection: Privacy settlement, proof procurement, and verified compute have different users, bottlenecks, collateral requirements, and political risks. Combining them in one token imports unrelated failures, cross-subsidizes weak services, and mistakes shared branding for monetary coherence.

Response:

- The objection is valid against any architecture that assumes one-token coherence. This thesis no longer assumes it.
- The Triad Coherence Test ([Section 7.3](#)) compares one native monetary asset, a neutral reserve plus service-specific credits, and shared settlement plus modular domain collateral.
- Architecture A survives only if common security and liquidity create more value than cross-domain governance and wrong-way risk destroy. If Architecture B or C performs better, the

thesis narrows to an open service stack beside an existing reserve asset.

31.1.4 “Utility does not imply money.”

Objection: Electricity, bandwidth, and cloud compute are all indispensable, but claims on them are not stores of value. Why should Privacy, Proofs, and Compute be different?

Response:

- The objection is correct as stated. Utility alone does not create money. The [Value Capture Lemma](#) establishes only whether service value accrues to the base asset through fees, burns, collateral, issuance discipline, and constrained bypass.
- Monetary treatment additionally requires stress-deliverable service, a persistent self-custodied loss-bearing holder constituency, and the holder-side service flow of [Section 11.10.8](#). If those fail, the system can be useful infrastructure and the base asset a service or cash-flow claim without being a store of value.
- The telemetry regime makes the distinction testable: the Value Capture Board reads service accrual, the Native Monetary Buyer Map reads risk absorption, and the Layer 0 board reads DVC.

31.1.5 “The asset can be bypassed.”

Objection: Users may need proofs, privacy, and compute, but they can buy them from AWS, a prover marketplace, a privacy wallet, or a stablecoin-based service without holding the native asset.

Response:

- The objection is correct unless the protocol enforces native value capture. The SoV thesis requires that core fees, collateral, slashing, priority access, settlement, or governance-critical operations be denominated in the native asset; that a material share of fees be burned or retired; and that operators lock the asset to provide service.
- If users can consume equivalent triad capacity without touching the asset, the system may be useful infrastructure, but the monetary thesis fails. This is Red Line 6 ([Section 28.6](#)).
- The honest answer is: this risk is real and must be continuously monitored. The Economic Coverage Board and Value Capture Board ([Chapter 24](#)) track native-asset fee share and bypass indicators.

31.1.6 “Gold still beats this because gold needs no network.”

Objection: Gold requires no software, no telemetry, no cryptographic agility, and no functioning power grid. Any network-dependent asset is structurally inferior as a store of value.

Response:

- Correct under some failure modes. That is why gold is a bridge and permanent non-digital

redundancy asset (Section 27.1, Section 27.1.1), not a discarded precursor.

- The triad is not “better than gold” universally; it is better suited to dense digital civilization if the network remains reachable and verifiable. The two are complements across different failure surfaces, not substitutes competing for the same use case.

31.2 Technical and Infrastructure Objections

The next group disputes whether the triad can be supplied at all under adversarial conditions—whether anyone wants it, whether useful-work markets can stay decentralized, whether proofs deliver what they appear to, and whether the physical substrate can be denied to us.

31.2.1 “Users don’t care about privacy or proofs.”

Objection: People trade convenience for privacy all the time; most don’t verify anything. Why build a SoV thesis on properties most users won’t touch?

Response:

- Users may not **ask** for privacy or proofs, but they **suffer** when they’re absent: identity theft, surveillance, misinformation, censorship, frozen funds.
- The stack is not asking end-users to run verifiers or design circuits; it’s asking:
 - infra operators to run verifiers,
 - institutions to demand receipts, and
 - developers to call PAL/PRK instead of opaque APIs.
- The “user” that cares most may be a treasury, DAO, insurer, or regulator—actors who must explain themselves.

If, after a decade of increasing repression and AI-mediated reality, nobody is willing to pay for privacy, proofs, or verified compute, then the triad doesn’t become money. The thesis is that the opposite is more likely.

31.2.2 “PoUW will centralize / can’t compete with hyperscalers.”

Objection: Useful-work mining sounds good, but hyperscalers and incumbents will always dominate; you’ll just rebuild cloud oligopolies inside a blockchain.

Response:

- Hyperscalers already dominate raw compute; the point of PoUW is not to out-compete them on price, but to:
 - turn **verified units of useful work** into a commodity;
 - ensure **anyone** can verify;
 - keep entry for provers open at the margin.
- Layer-0/4/6 design fights centralization by diversifying hardware profiles; measuring and

publishing prover concentration; using rewards and Work Credit policies to favor diversity; allowing hyperscalers to participate, but not to be **the only ones**.

If the market decides that centralized compute is always “good enough” and verifiability never matters, then AI Money stays a nice phrase. The thesis is that high-stakes actors—finance, defense, safety-critical infra—will demand **verifiable** compute from multiple vendors.

31.2.3 “Proofs do not prove truth.”

Objection: The thesis calls these instruments “Attestation Money” and speaks of proofs as if they settle disputes. But proofs only verify computation under stated assumptions. They don’t tell us whether the inputs were honest, the model was appropriate, or the conclusion is socially meaningful.

Response:

- Correct. This is why the thesis uses “attestation” rather than “truth.” Proofs do not create truth; they bound disputes. They make specific claims cheap to verify: origin, custody, computation, policy compliance, and settlement finality.
- The thesis is not that cryptography tells us what is true, but that it reduces the surface area over which institutions must be trusted. That reduction is valuable even if it is bounded.
- If the thesis overclaims—if it suggests that proofs solve epistemology rather than computation verification—the reader should substitute “bounded attestation” wherever “truth” appears. The economic argument does not require semantic truth; it requires cheap, public verification of specific claims.

31.2.4 “They’ll just shut off the internet / app stores.”

Objection: Sovereigns can simply block the network or remove apps from stores.

Response:

- Total, permanent shutdowns are blunt and politically costly; partial, targeted throttling is more likely. The stack assumes this and designs for it:
 - Multiple obfuscated transports (Layer 1).
 - Content-addressed updates and offline installers (Layer 2).
 - Sideload paths for clients.
- Treat reachability and update-health as SLOs (VerifyReach): if clients in censored regions can still fetch proofs and updates via at least one path, comms resistance is doing its job.

31.2.5 “This burns too much energy.”

Objection: PoUW is just PoW with extra steps; it still wastes energy.

Response:

- All monetary substrates consume scarce resources—geology, enforcement, balance-sheet ca-

capacity, or energy. The premise of the objection is that SHA-256 hashing is waste, and that premise is rejected here as firmly as the objection is: the energy buys unforgeable costliness under a work function nobody outside the system can price, and that is what the money is made of (Section 31.1.1). PoUW does not correct a defect. It binds the same class of expenditure to work that has an independent buyer.

- The right question is energy per verified FLOP or per proof unit, and whether that trend is improving. If energy-per-receipt falls while verified-capacity-per-token rises, the system is delivering more *verified capacity* per unit of energy. It is not thereby delivering more *monetary objectivity*: on that axis a work function with an external buyer is weaker, not stronger, and the comparison should not be run in a single direction and reported as a win.
- **Measurable answer:** Facility Energy Receipts (FERs) and “work per kWh” metrics make energy usage auditable. Energy efficiency shows up in VERIFYPRICE (cost component) and in Layer-0 telemetry.

31.2.6 “You will lose the competition for electrons.”

Objection: This is the serious version of the energy objection, and it is not about waste. Verification workloads will compete for power against industrial demand and state-backed AI build-outs. Those competitors have sovereign balance sheets, national-champion status, and priority in the interconnection queue. Third-party proving loses that contest, and no amount of telemetry changes the outcome.

Response:

- **This objection is substantially correct**, and the thesis now treats it as a threat class rather than a talking point (Section 5.5). Curtailment, tariff discrimination, interconnection denial, and load prioritization are real instruments, administratively cheap, and rarely legible as repression.
- **But it applies to proving, not to verification.** These are different exposures (Section 15.10.1). Verification runs on reference hardware at the edge in small amounts; its exposure is hardware obtainability, not bulk power. Losing the bulk-power contest raises the cost of *producing* proofs. It does not by itself break “anyone can verify,” which is the hinge.
- **Where it does bite, we measure it and let it bind.** Sovereign optionality (Section 15.8) aggregates curtailment exposure, fuel mix, backup depth, and jurisdictional risk into a published index; risk haircuts price fragility into issuance (Section 23.6.4); Red Line 13 retires the thesis if verification affordability becomes a sovereign policy variable.
- **The honest limit:** if a jurisdiction decides third-party proving will not be energized, cryptography does not overrule the grid. The response is dispersion and honest reporting. A thesis that claimed otherwise would be lying about physics.

31.3 State, Political, and Governance Objections

These objections concern power: whether states will tolerate lawful privacy, whether our own instruments can be turned against their users, whether a state-integrated stack simply outperforms an open one, and whether the thesis is analysis or preference. The most serious charge in this group is not that we lose, but that we win and become the thing we objected to.

31.3.1 “Governments will never allow lawful privacy at scale.”

Objection: Sovereigns will not tolerate unbreakable privacy and bearer-like digital money; they will regulate and block until these systems are marginal.

Response:

- Some governments will indeed oppose; others will see advantages in verifiable, receipt-backed compliance; lower settlement risk; reduced dependence on foreign platforms and currencies.
- The architecture is about **resilience**, not legal victory: it assumes partial repression; spreads hardware, comms, and governance across jurisdictions; keeps protocol-level custody and identity neutral.
- Lawful privacy is engineered, not begged for: viewing keys and receipts allow compliance without re-centralizing custody.
- **Critically:** The system designs *neutral infrastructure with consented disclosure*, not “dark finance.” Policy compliance is predicate-based (ZK proofs of allowlist membership, authorization, jurisdiction), not graph-inspection.

If all major jurisdictions converge on banning any form of non-custodial digital value, a lot more breaks than this stack. The design goal is: **if even a few open jurisdictions remain, and some gray-market paths exist, can the triad continue to function?**

31.3.2 “Receipts will become surveillance tools.”

Objection: All these receipts and proofs will just become a new surveillance layer. Governments will demand access; the system will comply; privacy dies.

Response:

- Receipts are **privacy-preserving by design**: they prove predicates (membership, compliance, authorization), not identities. A receipt proves “sender is in cleared set” without revealing which entry.
- Disclosure is **consented and scoped**: viewing keys reveal specific flows to specific auditors for specific time windows, not universal access. There is no “master key.”
- **Constitutional constraint**: “Policy = predicates, not graph inspection” is a hard constraint. Any policy requiring universal tracing is treated as incompatible with the monetary design.

- The system **explicitly rejects** policies that require global traceability. If a regulator demands “show me all transactions,” the answer is: “We can prove compliance with specific predicates; we cannot provide a surveillance feed, because the architecture doesn’t support it.”

If receipts become surveillance tools, the design has failed. The telemetry should make this visible: if most flows require real-name disclosure, the Settlement & Privacy Board will show it.

31.3.3 “This just gives the state better compliance tools.”

Objection: Selective disclosure, receipts, and viewing keys sound like privacy, but they are really a more efficient surveillance and compliance apparatus for the administrative state.

Response:

- The stack must distinguish selective disclosure from global graph inspection. Predicates (“sender is in cleared set,” “amount is within range”) reveal facts, not biographies. If proofs become dossier infrastructure, agency preservation fails—this is the ninth SoV requirement (Item 9) and is tracked directly by the social coercion metrics (Section 25.1.2) and the Agency Preservation Board (Section 24.9).
- The constitutional constraint in Section 25.1 (“policy = predicates, not graph inspection”) is a hard design boundary, not a marketing claim. Any policy requiring universal traceability is treated as incompatible with the monetary design.

31.3.4 “The closed sovereign stack simply wins.”

Objection: Chapter 30 concedes the case against itself. A state that controls energy, industry, compute, payments, and identity will out-build any open network, and will deliver better up-time and lower cost while doing it. Open verifiable systems are a luxury good for jurisdictions that have not yet had to get serious.

Response:

- **On capability, this may be right**, and the thesis does not contest it. A competent closed stack can plausibly beat an open one on cost, latency, uptime, and build speed. Directed capital does not wait for market clearing.
- **Capability was never the claim.** The thesis argues that certain properties—non-custodial settlement, private-by-default disclosure, portable identity, independently checkable receipts, practical exit—are what make an asset a store of value under repression (Chapter 4). A closed stack does not provide those at any level of capability, because withholding them is what makes it closed.
- **The two are not distinguishable by performance metrics**, which is exactly why the scoreboard must include agency and exit (Section 30.9). Identical VERIFYPRICE and uptime, opposite answers on whether a user can leave.

- **The real risk is not losing to it, but becoming it.** Red Lines 11 and 12 exist for this. An open stack that ends up permissioned, hyperscaler-hosted, and disclosure-mandated has not been outcompeted; it has converted. That is the failure mode worth monitoring, and it is monitored.

31.3.5 “A perfect ledger hardens an unequal starting state.”

Objection: Cryptographic verification solves alteration, double-spending, and provenance — then faithfully preserves whatever distribution existed when the snapshot was taken. If incumbents arrived with the assets and entrants must acquire them from wages at post-repricing valuations, a trustless ledger does not widen access; it makes the incumbents’ claims cheaper to hold and harder to contest. The stack could win technically and reproduce, more efficiently than fiat did, exactly the exhaustion dynamics of [Section 3.6](#) (“The Vibececession Was Real”; Sources).

Response:

- **The objection is correct about ledgers and has always been priced in.** Nothing in the thesis claims verification redistributes. The Value Capture Lemma’s conditions govern whether service value accrues to the asset; they are silent on who held the asset first, and a perfectly neutral ledger is neutral about incumbency too. Conceding this is not a retreat — the thesis never claimed otherwise.
- **But neutrality of record is not neutrality of access.** The fiat comparison is not distribution-neutral either: it is a system in which the mapping from discipline to advancement ([Section 3.6](#)) failed while every instrument read green. The relevant comparison for the open stack is whether it *broadens the acquisition channel* relative to that baseline — open-admission proving markets, fee-denominated entry at small denominations, collateral that is capacity-earned rather than capital-inherited, no institutional gatekeeper between a wage and a claim. These are architecture choices, they are already in the design (open admission, [Section 31.2.1](#); Tier C service markets), and they are what distinguishes “neutral about incumbency” from “rigged for incumbency.”
- **The stack must be measured on the capability layer, not the formal layer.** Formal access (a valid balance) can coexist with practical exclusion (no plausible path from wages to a meaningful claim). The distinction and its telemetry belong to the ninth requirement ([Item 9](#)) and the participation line; the Agency Preservation Board and the buyer-map’s cohort telemetry are where a widening acquisition gap would surface. A stack whose self-custodied cohort stops growing from income and grows only from inheritance has failed a requirement even if every SLO holds.
- **What would refute the response:** if, at material adoption, new-participant acquisition of the base asset from labor income stays negligible while incumbent concentration rises — visible in buyer-map cohorts and holder concentration — then the open stack’s acquisition

channel claim is false and the objection stands in its strong form. That is a measurement, not a value judgment, and it is on a board.

31.3.6 “This is political preference dressed as analysis.”

Objection: If both stacks minimize trust and both are hardened, the preference for the open one is ideology. A national stack delivering energy, payments, healthcare, and compute serves more people better than a fragmented network of cryptographic hobbyists. Calling one “sovereignty for the person” is branding.

Response:

- **The normative commitment is real and stated**, not smuggled: the thesis’s purpose is preserving practical agency under monetary instability and administrative integration ([Item 9](#)). It does not pretend to be value-free, and a document that claimed neutrality here would be less honest, not more.
- **But the analytical content is separable and falsifiable.** The claim is conditional: *if* soft guarantees weaken, *then* a bearer asset may earn premium only where service remains deliverable, value accrues without bypass, and holders bear loss—measurable through DVC, fees, burns, collateral, settlement health, buyer quality, and verification cost. [Section 28.6](#) lists fifteen ways to show it false. Ideology does not usually publish its own kill conditions.
- **The competence objection is conceded.** A closed stack may well deliver more material welfare to more people. The thesis’s response is not that this is untrue but that it is a different question from what happens to holders when the administrator’s incentives change and exit is unavailable.
- **What would actually refute us** is set out in [Section 28.6](#): if institutions restore broad credibility without increasing administrative control, if users retain practical exit despite integration, and if open proof systems fail to provide economical alternatives, the thesis is wrong regardless of anyone’s preferences.

31.3.7 “AI will be enclosed anyway.”

Objection: Hyperscalers have too much capital, data, and distribution advantage. Decentralized verified compute will never be more than a rounding error next to the walled gardens.

Response:

- The thesis does not require decentralized verified compute to be cheaper than hyperscalers ([Section 31.2.2](#)). It requires it to be available, verifiable, and censorship-resistant when those properties carry a premium—the same logic developed for AI Homestead vs. AI Enclosure in [Section 5.6](#).
- The Homestead Ratio and Enclosure Risk Flag make the concentration question falsifiable rather than rhetorical. If hyperscaler share of VERIFYPRICE-tracked capacity stays persis-

tently above threshold, Red Line 11 (Section 28.6) triggers and the thesis concedes the enclosure path is winning.

31.3.8 “Governance will just re-centralize.”

Objection: The stack is too complex; whoever runs it will become the new chokepoint.

Response:

- Complexity does not have to mean opaqueness. The governance layer is deliberately thin: decisions are about parameters and SLOs, not about picking winners.
- Multi-jurisdictional foundations, transparent config changes, and slashing rules keyed to public telemetry (VERIFYPRICE, decentralization stats, corridor health) make governance legible.
- If a network cannot show **who changed what, when, and in response to which metrics**, it is not an SoV candidate—however elegant its whitepaper.

31.4 Market-Structure and Co-option Objections

The final group is specific to the Market Realization Plane, and it is the group most likely to be dismissed by protocol designers as somebody else’s problem. It concerns what happens when the asset succeeds financially without succeeding monetarily—and in particular the tempting inference that a rising price settles the argument.

31.4.1 “The asset will be co-opted, not bypassed.”

Objection: Even if the protocol enforces native value capture, the financial system will co-opt the asset through ETFs, treasury vehicles, margin loans, and regulated custody. Most holders will never use self-custody, privacy rails, or verification tools. The asset becomes a speculative vehicle inside the existing financial system, not an independent monetary base.

Response:

- The objection describes a real and ongoing dynamic. The thesis acknowledges it as a variant of bypass risk: co-option routes demand through custodial instruments that do not exercise the fee-burn-collateral loop.
- The defense is protocol-level: if privacy, settlement, and proof capacity can only be accessed through native protocol interactions (not custodial wrappers), then custodial holders have exposure to price but not to the monetary properties that justify the price. Over time, this creates a divergence: custodial exposure without utility is speculation; protocol-native usage with utility is monetary demand.
- The telemetry regime must track this: native-asset fee share vs. custodial wrapper volume, protocol-native settlement vs. exchange-settled trades, and self-custody ratio vs. custodial

concentration. If custodial wrappers dominate and protocol-native usage stagnates, the co-option objection wins.

- The strongest counter-argument is that the protocol stack must make self-custody and protocol-native usage *easier and more rewarding* than custodial alternatives—not by ideology, but by delivering privacy, verification, and settlement that custodians cannot replicate.

31.4.2 “This will just become another ETF.”

Objection: Regulated wrappers, treasury companies, and ETFs will absorb most institutional demand. The asset will become a reference price for a financial product, not a monetary rail.

Response:

- It might. That is why the Wrapper Dominance Ratio exists (Section 11.13). Price exposure without protocol usage is not monetary adoption, and the thesis names this failure mode explicitly rather than hiding behind price action.
- The telemetry regime is designed to catch this early: if WDR rises while native fee share, private settlement, and collateral lockups stagnate, Red Line 9 (Section 28.6) triggers before the price signal would suggest anything is wrong.

31.4.3 “The price went up, so the monetary thesis is working.”

Objection: The asset has appreciated substantially and institutions are buying it. That settles the argument.

Response:

- It settles nothing. Price appreciation can be produced by return-decoupled wrapper allocations, leveraged rebalancing, dealer hedging, passive inclusion, treasury-company issuance, or trend following — none of which touch a fee, a burn, a proof, or a settlement (Corollary 11.1).
- Monetary adoption requires native fee demand, burns, collateral, private settlement, proof consumption, verified compute, and healthy stack telemetry. Price is neither necessary nor sufficient evidence for any of these.
- The uncomfortable symmetric claim also holds: a falling price is not falsification. The thesis is tested on the Value Capture Board (Section 24.9), and interpreted — never tested — through VerifyFlow (Section 24.8).

31.4.4 “Wrappers increase access, so co-option is harmless.”

Objection: ETFs and custodians expand the buyer base, deepen liquidity, and give institutions a legal path in. That is unambiguously good for the asset.

Response:

- Partly true, which is what makes it dangerous. Wrappers can improve liquidity and insti-

tutional holdability. They can also centralize custody, bypass native value capture, reduce self-custody, and manufacture holders who never exercise the monetary properties being capitalized (Section 5.6).

- The outcome is empirical, not ideological. The question is whether wrapper growth *seeds* native use or *substitutes* for it, and the Wrapper-Native Growth Gap answers it directly (Section 24.8).
- Note what this thesis does *not* claim: that every ETF is an adversary. Unlevered wrappers can be genuinely useful. The failure is not their existence; it is the inability to distinguish accessibility from adoption.

31.4.5 “Market mechanics are external, so protocol designers can ignore them.”

Objection: Wrappers, dealers, and fund flows are outside the protocol boundary. Designers should build the stack and let markets do whatever markets do.

Response:

- External wrappers reach back inside. They affect treasury behavior, collateral demand, governance concentration, liquidity-provider incentives, public understanding, security-budget expectations, regulatory pressure, and the price that the protocol’s own participants use to make decisions.
- The correct posture is narrow but firm: the protocol must not *manage* its price, and must nonetheless *instrument* the structures through which its monetary claims are represented (Section 11.14). Measurement is not intervention.
- Ignoring the plane does not make it inert. It makes it unmeasured, which is precisely the condition under which the compositional adversary (Section 5.2) does its work.

31.5 Completeness Objections

This last group is not an attack on moneyness. It is an attack on the claim that moneyness is enough. A duration-neutral, verifiable, bearer-capable asset can still leave the physical stack unbuilt.

31.5.1 “Who finances the reactor?”

Objection: Neutral money and verifiable settlement can preserve value. They cannot, by themselves, finance a nuclear plant whose construction takes ten years and whose productive life spans sixty. Gold has no coupon. Bitcoin has no underwriting desk. Proofs establish that a computation occurred. Privacy prevents informational capture. None of these converts present savings into long-lived productive capacity. AfterFiat has told savers to exit the bond, then left civilization with no way to build the plant. That is not a monetary design. It is a

savings technology with a stack diagram.

Response:

- **The objection is correct as a completeness claim, and it is conceded as one.** Duration-neutrality of the monetary object (Chapter 4) is a requirement of repression resistance. It is not a theory of capital formation. Section 3.3 distinguishes duration of the claim from duration of the project for exactly this reason. A post-fiat order that refuses credit altogether will not build Layer 0. A post-fiat order that cannot tell collateral from credit will recreate the opaque leverage and maturity mismatch this thesis is trying to escape.
- **The thesis is not the duration warehouse, and Layers 0–6 are not a pension.** Reserve collateral should remain duration-neutral, politically hard to pin, and bearer-capable. Duration finance should remain labeled credit: underwriting, covenants, maturity transformation, loss-bearing equity, and institutions that can sit through marks. Proofs can make construction progress, energy receipts, Facility Capacity Receipts, covenants, and restructuring states checkable. They cannot delete time risk, inflation risk, construction risk, political risk, or technological obsolescence. If a project cannot attract a loss-bearing holder of time, it does not get built—and that failure is not a monetary-design bug.
- **Treasuries currently smash the two functions together.** They are used as safe collateral, as a yield curve, as pension-matching assets, as repo, and as the mechanism that converts savings into public expenditure. Green’s long-end analysis shows the duration function weakening while the collateral function is still treated as intact. Replacing Treasury reserve demand with gold or Bitcoin would not automatically replace duration and credit. The world would still need instruments through which savers finance governments, infrastructure, industry, housing, and corporate investment. The design problem AfterFiat actually poses is how to *separate* those functions without recreating discretionary bailout structures.
- **China’s answer is command; the current American answer is a rule.** Directed banks can be compelled to warehouse the interval. Default enrollment into market-value-weighted bond indexes cannot. Neither is the open-stack answer. The open-stack answer is: money stays duration-neutral; credit stays labeled and auditable; governance decides who may write the default and who bears the loss when the rule fails (Chapter 23, Section 3.6).
- **What would make the objection fatal rather than clarifying.** If the thesis began treating Work Credits, proofs, or the native asset as a substitute for long-duration project finance—a coupon by another name, a construction bond wearing a monetary label—then duration-neutrality would have been undone and the bondholder kill box would have been re-entered through the protocol door. That is a use this document refuses.

31.5.2 The Open Duration Warehouse

“Credit must remain labeled credit” is a boundary, not a financing architecture. An open stack therefore permits explicit duration-bearing instruments beside the duration-neutral base asset:

- proof-audited project notes with stated maturity, covenants, default states, and loss priority;
- revenue-linked capacity bonds whose payments depend on attested delivered service rather than promised nameplate output;
- power-purchase and compute-purchase agreements that contract defined quantities, locations, delivery windows, and force-majeure states;
- senior, subordinated, and first-loss infrastructure tranches with disclosed attachment and exhaustion points;
- pre-purchased capacity contracts that transfer delivery and basis risk explicitly;
- separately capitalized first-loss or insurance facilities with published exclusions and claims waterfalls;
- bankruptcy-remote project entities that isolate construction and operating claims from the base-asset constitution; and
- continuously attested cash-flow, collateral, construction-milestone, and restructuring waterfalls.

Proofs can reduce information asymmetry, attest milestones, audit revenue, and make covenant breaches public. They cannot eliminate default, commodity risk, construction delay, technology obsolescence, political seizure, or the need for loss-bearing capital.

Constitutional Separation of Money and Project Credit

No par guarantee, redemption promise, liquidity backstop, bailout, or emergency support for project notes, capacity bonds, purchase agreements, tranches, or insurance facilities may be transferred—explicitly or implicitly—to the base asset. Losses remain with the contractual credit and equity waterfall. Governance may verify and resolve the claim; it may not monetize it through the bearer asset.

This architecture does not promise cheap capital. It makes the duration holder, default boundary, and loss waterfall legible. If an infrastructure project cannot attract an investor willing to bear its time risk on those terms, the missing object is credit appetite, not money.

Chapter 32

Why These Could Become Money

It is fashionable to say that money is a shared hallucination. That line flatters our cleverness while excusing our passivity. Hallucinations cannot settle debts across adversarial jurisdictions; hallucinations do not finance supply chains. Money works because it is backed by a machine—sometimes a literal machine of war, sometimes a machine of law, and now verification machines that continuously produce scarce, indispensable utilities.

Post-Bretton-Woods, the machine was compliance. The next machine is verification. Cypherpunks did not abolish trust; they automated it. When privacy, proofs, and compute clear across neutral rails, money stops asking for favors and starts paying for facts.

We can restate the thesis as a conditional:

*If a dense digital civilization continues to rely on AI, global networks, and programmable markets, and if states continue to use repression and narrative control rather than explicit default, and if the full service path remains stress-deliverable, value accrues without bypass, and a self-custodied loss-bearing holder constituency persists, then a bearer base asset **may earn monetary premium** by preserving settlement, proof, compute access, and exit when ordinary substitutes weaken—provided project credit remains separate and without assuming the triad must share one token.*

The argument, restated as a conclusion chain:

1. Digital civilization structurally needs privacy, proofs, and compute.
2. These capacities can be standardized and verified through the seven-layer stack.
3. Standardized work can produce cheap-to-check receipts.
4. Receipts represent capacity only where the full path delivers usable service under stress.
5. Stress-deliverable receipts can support markets.
6. Markets can produce recurring fee flows.
7. Fee flows support a scarce base asset only under non-bypassable accrual rules.
8. A persistent, self-custodied, loss-bearing, regime-responsive constituency must warehouse that asset's risk.
9. Only then, if the asset remains liquid, neutral, verifiable, legally holdable, and agency-preserving, may it earn a store-of-value premium.

The asset may first behave like infrastructure equity—a claim on fee flows from useful services. Only later, if liquidity deepens, neutrality is sustained, and holdability matures across jurisdictions, does it potentially cross the threshold from “productive asset” to “monetary collateral.” That transition is not guaranteed; it is the thesis’s central conditional claim, and it has a base-rate problem the thesis should state rather than bury: **no asset in monetary history is known to have made this crossing.** Warehouse receipts, bills of lading, and standardized commodity claims — the closest historical analogs to Work Credits and fee units — acquired price and liquidity without ever acquiring monetary premium. The thesis is not describing a developmental stage other assets have passed through; it is proposing that a regime now exists in which the crossing could first occur. What follows is the argument for that possibility, not a precedent.

The media-transition literature offers a reading of that failure record, and it is offered as a reading, not as evidence: a printed scroll is still a scroll. Each of those precedent systems delivered the *content* of a new capacity—storage, rendering, compute—while keeping the *form* of the old one: a service purchased from an operator, priced as a service, held as an operating claim. Content without structure is exactly what a medium looks like in its first generation, and the utility-token record is that first generation’s ledger. On this reading the crossing the thesis proposes is not an increment on those systems but a change of kind—fees are the content; the holder-side stress-deliverable service is the medium—which is why the base rate is empty and why the empty base rate is not dispositive. It is also why the reading earns nothing: no link of [Section 7.1](#) strengthens because a pattern from another century rhymes. The chain still has to prove itself, link by link, in telemetry. The honest summary of the two moves: the empty base rate is *evidence against* — it raises the prior the crossing fails, as fifteen empty precedent rows should — and the change-of-kind argument is *the reason the evidence is not dispositive* — it identifies a property the precedents lacked, which changes what the base rate is a base rate of. Neither move makes the crossing likely; together they make it testable, which is what the red lines are for.

32.1 As a Store of Value

From the SoV lens, a credible store of value must be credibly scarce, cheap to verify in public, resist censorship and capture, have native demand that is not purely narrative, and avoid being a duration instrument whose real return can be pinned negative by policy.

Privacy, Proofs, and Compute may support a base asset that meets that brief, but the services themselves and their Work Credits are not presumed monetary:

- **Privacy** is purchased because some people and institutions must pay without broadcasting their graph: dissidents, NGOs, treasuries under capital controls, enterprises with sensitive payroll and vendor relationships. In a repression-heavy world, privacy is not a luxury good;

it is the hull that keeps savings from becoming an option owned by someone else.

- **Proofs** are purchased because “seeing is believing” has failed. Deepfakes, platform curation, and the liar’s dividend make any unproven artifact suspect. Regulated AI and finance regimes require auditable provenance and computation. In that world, proofs are not a niche; they are the affidavit layer of the digital order.
- **Compute** is purchased because intelligence is now a first-class input to production. FLOPs for training, inference, and proving are line items in budgets. Once those FLOPs are wrapped in proofs and standardized as canonical workloads, verified compute becomes a commodity that can be priced, hedged, and stored.

In a world that may choose **stealth default**—negative real yields aided by regulation—over explicit default, a durable SoV must be duration-neutral, peg-proof, stress-deliverable, and held by actors able to bear loss. Recurring utility can support value accrual; it cannot supply moneyness by itself.

Duration-neutrality of the *claim* is not, however, a theory of how plants get built. [Section 3.3](#) and [Section 31.5.1](#) state the completeness condition the rest of this chapter does not repeat: reserve collateral and duration finance are different functions; proofs can audit the second; they must not become it. A stack that supplies Privacy, Proofs, and Compute still requires some institution, outside Layers 0–6, capable of warehousing the construction interval. That institution is not the token.

32.2 As Stack

Seen from the stack angle, the triad is backed not by a metaphor but by a supply chain:

- Layer 0 keeps **machines** honest and powered (verifiable hardware, FERs).
- Layer 1 keeps **packets** flowing under censorship (VerifyReach).
- Layer 2 keeps **code** moving even when app stores and CDNs are hostile.
- Layer 3 keeps **identity** accountable without doxxing.
- Layers 4–5 make **work** and **value** flow through proofs and privacy rails.
- Layer 6 keeps **governance** and **telemetry** legible.

Triad instruments are claims with different seniority and risk. Work Credits represent typed future access or service capacity; FCR, FER, and PIDL artifacts provide evidence; project notes bear duration; LP and staking positions are derivatives or operating claims. Only the base asset is tested as a monetary candidate.

32.3 As Telemetry

VERIFYPRICE, VerifyReach, VERIFYSETTLE, and decentralization metrics are the constitution. They keep “trustlessness” from decaying into “trust the custodians.”

A system that cannot show its own health cannot be money; it can only be marketing.

32.4 Trading One Base Reality for Another

What changes and what stays the same?

What changes: The base reality that backs money. Instead of “the sovereign will repay” or “gold is scarce because geology,” the base reality becomes “these capacities are scarce, necessary, and verifiable.”

What stays the same: Money is still a claim on work. The work just becomes specific and measurable: proofs that anyone can check, privacy that anyone can use, compute that anyone can verify.

The operative word is **agency**. Monetary arrangements can be read as answers to the question: who has agency over their own savings, payments, and economic identity? Gold answers it with physical possession. Fiat answers it by delegation to institutions. Bitcoin answers it with self-custody and cryptographic keys, and answers it well enough that this thesis builds its settlement layer on top of the answer rather than around it. The triad addresses adjacent surfaces that self-custody does not reach: agency over verification (proofs), over private settlement (paying without being watched), and over productive capacity (compute without permission). Different surfaces, not a later stage ([Section 31.1.1](#)). A store-of-value premium, where it exists, accrues to assets that preserve agency on whichever surface is under attack.

See [Sections 32.2](#) and [32.3](#) for how the same claim reads through the stack and telemetry lenses.

Chapter 33

Conclusion: A Bell Labs for Privacy, Proofs, and Compute

The original Bell Labs turned information theory into cables, switches, and semiconductors that quietly reshaped the world. Nobody needed to know the math; they just made phone calls.

The mandate here is similar, but for a different substrate:

Turn privacy primitives, zero-knowledge, and verifiable compute into everyday infrastructure—receipts, rails, and verifiable machines—and test whether a separate bearer base asset can earn monetary premium by preserving agency through that stack.

That mandate breaks into three concrete programs:

1. **Monetary program.** Test the base asset against the nine-link chain and Native Monetary Buyer Map. Design Work Credits separately as typed, DVC-bounded service or capacity claims; do not ask them to warehouse monetary or project-duration risk.
2. **Stack program.** Build and harden Layers 0–6: verifiable machines, resilient comms, pseudonymous identity, PoUW and proof factories, privacy corridors, governance and telemetry.
3. **Telemetry program.** Keep VERIFYPRICE, VerifyReach, VERIFYSETTLE, and decentralization metrics live, public, and blunt. Treat dashboards as constitution.

33.1 The Quiet Mic Drop

If it does work, the moment of success will not look like an ICO or a conquest. The mic drop is quiet.

There will just be a decade in which:

- auditors ask for **receipts**, not screenshots;
- regulators accept **cryptographic attestations** as first-class evidence;
- enterprises budget for **proofs and privacy** the way they once budgeted for bandwidth;
- wallets and apps call `verify(receipt)` as casually as they call `https://`;

- and allocators treat triad assets not as exotic bets, but as part of the boring hedge against repression.

At that point, the thesis will have stopped being a thesis. It will have collapsed into plumbing.

33.2 Final Thoughts

Bitcoin established that **digital scarcity alone** can be money. That result stands, this thesis does not revise it, and nothing here improves on the mechanism that produced it (Section 31.1.1). The question taken up in these pages is a different one: whether **native utility that the world must keep buying**, secured by proofs anyone can check, can also carry monetary premium—or whether it only ever prices as a service.

Privacy will preserve agency.

Proofs will bound what can be faked—not create truth, but shrink the territory where forgery pays.

Compute will power intelligence—verifiable and honest, not merely performed.

When these three clear across neutral settlement, money stops being an article of faith and becomes what it always wanted to be: **a record of work that cannot be faked and does not need permission to move.**

In the decade ahead we will learn whether digital necessities, wrapped in proofs anyone can verify, can become candidate money for a world where verification is more valuable than permission.

33.3 Two Loops, Not One

One discipline before closing, because it is the discipline most likely to be abandoned in a bull market.

The Create/Compute → Prove → Settle → Verify loop tells us whether the cypherpunk monetary stack works. A second loop — Narrative → Wrapper → Allocate/Lever → Hedge → Price → Narrative — tells us how financial markets represent it.

The first loop turns indispensable digital capacity into native monetary value. The second can recognize that value, ignore it, amplify it, or counterfeit its appearance. A serious store-of-value thesis must instrument both. Without VERIFYPRICE, VERIFYREACH, and VERIFYSETTLE, we cannot know whether the stack is real. Without VerifyFlow (Section 24.8), we cannot know whether market price reflects native adoption or external mechanical demand.

Gold can be wrapped without ceasing to be gold. Bitcoin can be held in an ETF without ceasing to settle. But the more monetary use migrates from bearer ownership into custodial and synthetic claims, the more the financial representation can diverge from the very proper-

ties that justified the premium. The destination is not merely a higher price. It is a monetary object whose price, utility, custody, and settlement remain connected closely enough that each can still discipline the others.

VERIFYPRICE, VERIFYREACH, and VERIFYSETTLE verify the money. VerifyFlow verifies the story the market is telling about it.

33.4 Equip, Don't Manage

The purpose of the stack is not to make people more legible to institutions. It is to preserve the capacity to act when institutions become coercive, synthetic, or brittle. A CBDC can manage citizens. A custodial stablecoin can manage customers. A platform AI can manage users. A cypherpunk monetary stack must do the opposite: **equip actors**.

Privacy equips people to move value without becoming visible by default. Proofs equip people to verify claims without trusting platforms. Compute equips people and agents to purchase machine work without asking permission from hyperscalers. Telemetry equips the public to see drift before the stack becomes another soft guarantee.

Gold condenses geology. Bitcoin condenses thermodynamics. The triad condenses **verifiable agency**: the ability to pay, prove, and compute under adversarial conditions. Three hard facts, held in the present tense, none of them retiring another—and the third is the only one still under test.

This closes the circle opened in [Section 3.6](#). Fiscal dominance makes financial repression attractive; AI makes truth and cognition platform-mediated; household non-participation makes administrative management politically tempting; and physical bottlenecks make compute, power, and hardware strategically scarce. Gold, silver, Bitcoin, and hard assets remain the bridge ([Sections 27.1](#) and [27.1.1](#)). Privacy, Proofs, and Compute can support an associated base-asset monetary candidate only when the full nine-link chain proves, instrument by instrument, that the service stays deliverable when it is attacked, that demand cannot route around the asset without paying for it, and that the people holding it can bear the loss. Every instrument named in this thesis exists to make one of those three sentences measurable.

And the circle has a household edge the sovereign layers do not supply. Proof tells us whether a claim is true; representation tells us whom the claim describes; capability tells us what participants can actually do as a result. A stack that verifies everything and represents almost no one is a more efficient ledger for the same exhaustion dynamics it was built to escape, which is why the representation discipline travels with every aggregate this thesis publishes and the capability layer travels with the ninth requirement. Equipping actors, in the end, means more than non-custodial keys: it means the architecture does not mistake extraction for resilience because it measured itself in dollars rather than people.

The thesis fails if the dashboards fail, if verification centralizes, if privacy rails cannot stay lawful and usable, or if value does not accrue to the asset. Those are not escape clauses; they are the terms of the bet.

The question this stack answers is not whether its ledgers balance. It is whether the people inside them are owners or tenants.

Sources

Theory & Framing

Davidson, J. D., & Rees-Mogg, W. (1997). *The Sovereign Individual: How to Survive and Thrive During the Collapse of the Welfare State*. New York: Simon & Schuster / Touchstone.

Booth, J. (2020). *The Price of Tomorrow: Why Deflation is the Key to an Abundant Future*. Stanley Press.

McLuhan, M. (1964). *Understanding Media: The Extensions of Man*. New York: McGraw-Hill. Cited for framing only: that a medium's structural effects arrive as externalities of adoption rather than as features anyone selects, and that the content of a new medium is initially the old medium's content. Used in [Section 2.0.1](#) and the Executive Memo as pattern, not as evidence; no claim of this thesis depends on it, and nothing in it is offered as measurement.

McLuhan, M. (1962). *The Gutenberg Galaxy: The Making of Typographic Man*. Toronto: University of Toronto Press. Cited for the sensorium framing—that a medium restructures the ratio of the senses before anyone notices—and for print's restructuring of the private reader. Same limits as above: synthesis, not experiment; mechanism, not quantities.

Eisenstein, E. L. (1979). *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early-Modern Europe*. Cambridge: Cambridge University Press. Cited in [Section 2.0.1](#) for the standardization thesis: that typographic uniformity preceded and enabled the standardization of physical production, and that print's consequences were causal and profound. It is cited for mechanism only. The thesis does not adopt any necessity claim—that the industrial revolution was possible only because of movable type—and no argument here rests on one.

Ong, W. J. (1982). *Orality and Literacy: The Technologizing of the Word*. London: Methuen. Cited for interiority: that sustained silent reading is a technology-dependent practice and the basis of the private self as an epistemic unit. Used in [Chapter 8](#) to place private settlement in a longer arc. Framing only; the thesis's claim about payments as the load-bearing privacy surface is unchanged by it.

Anderson, B. (1983). *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. London: Verso. Cited for one connection only: that print created mass simultaneous identity, and that the Native Monetary Buyer Map is how a monetary constituency renders itself legible to itself. The analogy is structural; no claim about adoption is derived from it.

Culkin, J. M., S.J. (1967). A Schoolman's Guide to Marshall McLuhan. *America*, 116(9), 232–235. Cited for attribution: “We become what we behold. We shape our tools and thereafter they shape

us” is Culkin’s line, routinely misattributed to McLuhan. It appears in this thesis, if at all, as Culkin’s.

Kaldor, N. (1939). Speculation and Economic Stability. *The Review of Economic Studies*, 7(1), 1–27. Cited in [Section 11.10.8](#) for one general mechanism only: that a service flow accruing to the holder of a physical good, rather than to the holder of a claim on it, can carry a price a cash-flow model does not reproduce. It is not cited for gold, which is a near-zero-convenience-yield commodity whose price residual is conventionally attributed to monetary and portfolio demand rather than to storage convenience; and it is not a claim that a triad asset is a storage commodity. No magnitude in this thesis derives from it.

Working, H. (1949). The Theory of Price of Storage. *The American Economic Review*, 39(6), 1254–1262. Companion to Kaldor, cited for the same mechanism and subject to the same limits: the storage-theory account of convenience yield in consumption commodities, where the service flow is an option against stockout. The analogy in [Section 11.10.8](#) is to the mechanism, not to the setting.

Market Structure, Flows & Inelastic Markets

Green, M. W. (2026, July 26). A Semi-Theory of Almost Everything: Who Was Actually Setting Semiconductor Prices This Spring? Tier1 Alpha. *The flow-elasticity, holder-recycling, and leveraged-rebalancing machinery in [Appendix H](#) is adapted from this work; the generalization to monetary telemetry is ours, as are any errors in it.*

Green, M. W. *Anchors Aweigh, My Boys.* The automatic-versus-marginal buyer distinction, the observation that market-value-weighted bond indexes absorb less DV01 per dollar after a selloff, and the claim that the quality of long-duration demand has deteriorated are drawn from this work and used in [Section 3.1.2](#), [Section 3.3](#), [Section 24.8.7](#), and [Section 24.8.8](#). Commentary rather than peer-reviewed research; used for mechanism, not for quantities. Tactical claims and the attribution of any particular yield move uniquely to duration withdrawal are listed and excluded in [Appendix I.4](#).

Green, M. W. *What the Treasury Needs.* Used selectively for the distinction among market repair, fiscal management, financial repression, and administrative enclosure, and for the proposition that policy can change duration-bearing capacity without erasing economic loss. Commentary rather than peer-reviewed research; used for mechanism, not for the proposal’s illustrative quantities or for a claim that any facility is necessary.

Green, M. W. *Sometimes, You Get What You Need.* Used as a scenario-framing example for the market-repair branch in [Section 27.2](#). Its tactical sequence for rates, currencies, gold, mortgages, and equities is quarantined rather than adopted as forecast.

Green, M. W. *The Vibecession Was Real.* Used for mechanism, not quantities: the dollar-weighted versus person-weighted aggregation distinction and its covariance identity (adopted as the representation discipline on board aggregates, [Section 24.8.8](#)); the price-level-versus-inflation-process-versus-

balance-sheet-hysteresis decomposition and the consumption-through-depletion mechanism (adopted in the household-transmission observables of Section 27.2 and the exhaustion gradient of Section 3.6); the state-versus-household resilience distinction $R_{\text{state}} \neq R_{\text{household}}$ and the formal operating-surplus/resilience-stock model (adopted in Section 30.9); and the lived-solvency framing behind the incumbency objection (Section 31.3.5). Commentary rather than peer-reviewed research. Excluded: the essay's illustrative simulation percentages (presented there as mathematical possibility, not prevalence), its pricing-power magnitudes (anecdotal), any psychological attribution to named officials, and its secondary statistics, which are treated as provisional pending tracing to underlying series.

Gabaix, X., & Koijen, R. S. J. (2021). In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. NBER Working Paper No. 28967. *Cited for the inelastic-markets premise that flows move prices at all, not for the market-impact function used in Appendix H; that square-root form comes from the market-microstructure execution-impact literature.*

Cheng, M., & Madhavan, A. (2009). The Dynamics of Leveraged and Inverse Exchange-Traded Funds. *Journal of Investment Management*, 7(4).

Energy, Industrial Capacity & Sovereign Optionality

Doomberg. *Fire Horse: analysis of Chinese energy strategy, coal, refining, conversion capacity, and stockpiling. Used in Chapter 30 to characterize a strategy of conversion optionality and “resilience over efficiency.” This is an analyst presentation, not peer-reviewed work; figures cited are its accounting rather than independent measurement, and no claim in this thesis is load-bearing on them.*

Doomberg. *Endangered Specious. Used selectively for the mechanism that gross feedstock is not deliverable output when conversion capacity is the binding cut. Commentary rather than peer-reviewed research; used to motivate the Topological Scarcity Lemma and DVC, not for article-specific quantities or a conclusion that aggregate abundance guarantees local resilience.*

Doomberg. *A Minor Detail. Used selectively as a stress scenario illustrating common-cause dependence across fuel, logistics, cooling water, and generation. The geopolitical sequence is not treated as a base case or tactical forecast, and no unsourced quantity enters the model.*

Doomberg. *Over Target. Used selectively for the principal-agent mechanism in the Sponsor Dependency Board: a subsidized operator and its sponsor may have different terminal objectives. Claims about the intent behind any particular attack or escalation path are not adopted.*

Gromen, L. *The Forest for the Trees, and market commentary of August 2026. Three framings are drawn from this work: sovereign debt as a claim on future real surplus, with monetary credibility jointly determined by productive capacity and the state's ability to mobilize it (Section 30.1); the reflexive dependence of fiscal capacity on collateral values (Section 3.2); and the reshoring trilemma (Section 30.3).*

Commentary rather than peer-reviewed research; used for framing, not for quantities. Quantitative claims from the same source that could not be verified are listed and excluded in [Appendix I.4](#).

Note on the synthesis. The layered reading in [Chapter 30](#)—physical throughput beneath sovereign balance sheet beneath market realization beneath verification—is our construction, assembled from the above together with the market-structure sources in the preceding section. Sovereign Optionality adapts pathway diversity to facility and network scale; Delivered Verified Capacity restores the conversion-throughput term as scenario-specific service max flow. The Topological Scarcity Lemma, Pressure–Capacity Corridor, and their application to AfterFiat are our synthesis. Errors in that adaptation are ours.

Bessent, S. (2026, June 23). *American Economic Statecraft in the 21st Century*. Remarks at the Economic Club of New York’s America 250 Gala Dinner. U.S. Department of the Treasury, press release sb0539. *A primary source, cited in [Section 30.3](#) for stated doctrine only. Quotations are verbatim. A policy speech is evidence of intent and framing, not of outcomes, and this document draws no inference from it about what will actually be built.*

Kimi Team. (2026). *Kimi K3: Open Frontier Intelligence*. arXiv:2607.24653. *Cited in [Section 30.4](#) for a single narrow proposition: that openly released weights have approached the proprietary frontier closely enough to matter economically. The paper reports the model trailing the strongest proprietary systems overall, and that qualification travels with the citation. No claim about inference pricing is drawn from it.*

Macro: Debt, Financial Repression, Real Yields

Gaspar, V., Gonçalves, C. E., & Poplawski-Ribeiro, M. (2025, September 17). Global Debt Remains Above 235% of World GDP. *IMF Blog / Global Debt Monitor*. International Monetary Fund.

International Monetary Fund. (2025). *Global Debt Monitor 2025*. IMF Global Debt Database.

International Monetary Fund. (2026, April). *Fiscal Monitor, April 2026*. IMF. [Reports global public debt at just under 94% of GDP in 2025, projected to reach 100% by 2029.]

Reinhart, C. M., & Sbrancia, M. B. (2015). The Liquidation of Government Debt. *IMF Working Paper 15/7*. International Monetary Fund.

Board of Governors of the Federal Reserve System. (2025, October 15). The Cross-Border Trail of the Treasury Basis Trade. *FEDS Notes*. *Cited in [Section 3.1.2](#) for the scale and financing of Cayman-domiciled hedge-fund Treasury positions and for the resulting mismeasurement in the TIC data and the Financial Accounts. Staff research, not a Board policy position. The note itself cautions that the reported gap is not attributable solely to the cash–futures basis trade, and that caution is reproduced in the main text.*

Rose, J. (2021). Yield Curve Control in the United States, 1942 to 1951. *Economic Perspectives*, Federal Reserve Bank of Chicago, 45(2).

Social / Disinformation / Deepfakes / Provenance

“Twitter Files.” (2022–2023). Series of internal document releases published by independent journalists via the Twitter/X platform. [Primary documents; this thesis cites the releases as evidence of government–platform coordination, not as dispositive proof of any particular claim.]

U.S. House Committee on the Judiciary (Republican Staff). (2023). *The Weaponization of CISA: How a “Cybersecurity” Agency Colluded with Big Tech and “Disinformation” Partners to Censor Americans*. Staff Report.

U.S. Department of State. (2021–2025). Global Engagement Center (GEC): About Us. U.S. Department of State.

Harwell, D., & Patton, D. (2025, October 22). We uploaded a fake video to 8 social apps. Only one told users it wasn’t real. *The Washington Post*.

UNESCO. (2025). *Deepfakes and the Crisis of Knowing*. UNESCO briefing.

Ahmed, S., et al. (2024). Social Media News Use Amplifies the Illusory Truth Effects of Deepfakes. *Journalism & Mass Communication Quarterly*.

Policy Clock: EU AI Act

European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union.

Verification Asymmetry & PoUW (MatMul)

Komargodski, I., Lin, S., & Weinstein, O. (2025). Proofs of Useful Work from Arbitrary Matrix Multiplication. *arXiv:2504.09971*. Submitted April 2025; last revised November 2025. [Note: the optimal-security claim is described as a *conjecture* in the abstract, not a proven guarantee. The paper supports “promising research direction,” not “settled deployment.”]

ZKML / Verifiable ML

Peng, Z., Wang, T., Zhao, C., et al. (2025). A Survey of Zero-Knowledge Proof Based Verifiable Machine Learning. *arXiv:2502.18535*.

FHE / FHE Rollups / OpenFHE / Fhenix

Zyskind, G., Erez, Y., Langer, T., Grossman, I., & Bondarevsky, L. (2024). FHE-Rollups: Scaling Confidential Smart Contracts on Ethereum and Beyond. *Proceedings of BSCI '24*. ACM.

OpenFHE Project. (2025). OpenFHE v1.4.0 – Development Release Announcement. GitHub.

Fhenix. (2023). Introducing FHE-Rollups: Scaling Confidential Smart Contracts on Ethereum and Beyond. Fhenix blog.

Bitcoin Transport Privacy (BIP-324, Core 27.0)

Bitcoin Core Developers. (2024, April 16). Bitcoin Core Version 27.0 Released. Bitcoin-Core.org.

Mehta, D., Ruffing, T., Schnelli, J., & Wuille, P. (2023). BIP-324: Version 2 P2P Encrypted Transport Protocol. Bitcoin Improvement Proposals.

Lightning Receiver Privacy (BOLT12 Offers)

No Bullshit Bitcoin. (2024, July 4). Phoenix Wallet v2.3.1 & Phoenixd v0.2.0: BOLT12 Offers. NoBSBitcoin.com.

No Bullshit Bitcoin. (2024, September 24). BOLT12: Offers Officially Merged into Lightning Specification.

Ecash / Federations (Fedimint / Fedi)

Bitcoin Magazine. (2025, October 29). From Stealth to Scale: Fedi Unveils Multi-Sig Guardians for Federated Bitcoin Ecash Mints.

Fedimint Project. (2024–2025). What Does the Federation Do? Fedimint documentation.

Privacy Rails: Zcash, BTC↔XMR Swaps

Electric Coin Co. / Zcash Foundation. (2024). *Zcash Protocol Specification – NU6 Activation and Second Halving*. Zcash Protocol Specification, Rev. 2025.6.0-105.

ReitwieSSner, C. et al. (COMIT Network). (2020). Monero–Bitcoin Atomic Swaps. COMIT blog and GitHub.

UnstoppableSwap / eigenwallet. (2025). UnstoppableSwap GUI / eigenwallet: Non-custodial XMR↔BTC Atomic Swaps. GitHub.

Zcash Roadmap / NU7 Discussions

Zcash Community Forum. (2024–2025). Accelerate NU7: Let's Deliver a Smaller Network Upgrade Sooner. Governance thread.

Ethereum Privacy L2 (Aztec)

Steinrueck, D., & Sammara, A. (2025, May 2). Testnet Retro – Launch Day: Aztec Public Testnet Went Live. Aztec Network blog.

Coindesk. (2025, May 1). “Everything Is Encrypted”: Aztec's Privacy Rollup Hits Testnet Amid Growing Demand.

IBC-Wide Shielding (Namada)

Namada Team. (2025, June 20). Namada Mainnet Launch Is Complete! Namada blog.

Namada Team. (2023, January 31). Understanding the MASP and CC Circuits. Namada blog.

Proof Markets / PoUW / Privacy-First L1s

Aleo. (2025). Consensus and Proof of Succinct Work (PoSW). Aleo developer documentation.

RISC Zero / Boundless. (2025). Boundless: A Universal ZK Compute Layer. Binance Research.

Succinct Labs. (2025, August 5). Succinct Prover Network Mainnet Launch (PROVE).

Nockchain Foundation. (2025). Nockchain Is Now Open Source; Nockchain Issuance Schedule and Dumbnet Launch. Nockchain.org.

Psy Protocol (formerly QED). (2025). Psy: A Scalable, Privacy-First L1 Secured by Proof-of-Useful-Work 2.0.

Bittensor. (2025). Yuma Consensus. Bittensor Docs.

Ambient. (2024). Ambient Litepaper v1: Proof of Logits and Consensus.

Transport, Privacy, SDKs & Related Infrastructure

No Bullshit Bitcoin. (2024, July 11). The New Phoenix: A 3rd Generation Lightning Wallet.

Fedimint. (2025). Fedimint Lightning Gateway Uses LDK Node to Simplify Deployment and Liquidity Management. Lightning Dev Kit blog.

Ethereum Privacy Roadmaps (PSE / Privacy Cluster)

Privacy Stewards of Ethereum (PSE). (2025, September 12). PSE Roadmap: 2025 and Beyond. PSE.dev blog.

Ethereum Foundation. (2025, October 8). The Ethereum Foundation's Commitment to Privacy. Ethereum.org blog.

CoinMarketCap Academy. (2025, October 9). Ethereum Foundation Launches Privacy Cluster for Enhanced Blockchain Security.

Vitalik on Openness & Verifiable Hardware

Buterin, V. (2025, September 24). The Importance of Full-Stack Openness and Verifiability. Personal blog (vitalik.eth.limo).

C2PA / Provenance (Supporting)

Creative Commons / Coalition for Content Provenance and Authenticity (C2PA). Content Credentials Standard Documentation.

Software Update Security & Transparency (Layer 2 Prior Art)

The Update Framework (TUF). (2011–). *The Update Framework: A Framework for Securing Software Update Systems.* specification and reference implementations, theupdateframework.io. *The supply-chain invariants I1–I5 in Chapter 17 (threshold signatures, freshness/anti-rollback, reproducibility, compromise response, multi-path availability) restate, in protocol-specific form, guarantees TUF established for package repositories; Uptane (Kuppusamy et al., 2016–present) adapts the same framework to automotive update streams.*

Laurie, B., Langley, A., & Kasper, E. (2013). Certificate Transparency. RFC 6962, IETF. *The append-only, publicly auditable release log of Section 17.3 is the CT construction applied to software artifacts; witness cosigning in the style of Sigsum (sigsum.org) is the lightweight variant.*

Reproducible Builds Project. (2014–). *Reproducible Builds: Setting the Record Straight* and associated specifications. reproducible-builds.org.

Censorship Measurement (Layer 1 Prior Art)

Open Observatory of Network Interference (OONI). (2012–). *OONI Probe and Measurement Methodology*. ooni.org. *The VerifyReach measurement design (Chapter 16) — volunteer vantage points, canary requests, geographic/ASN stratification, public aggregation with p50/p95 reporting — follows the measurement methodology OONI has operated at scale for over a decade; the divergence triggers and WC-incentivized vantage pool are the extensions.*

Dingledine, R., Mathewson, N., & Syverson, P. (2004). Tor: The Second-Generation Onion Router. *Proceedings of USENIX Security '04*. *Tor and I2P as first-class transports (Chapter 16) are prior deployments of the transport-agility and receiver-privacy principles; pluggable transports are the obfuscation layer referenced by “obfs-class shims.”*

Sourcing Note

This thesis cites primary documents, academic research, official specifications, and direct industry publications where possible. Where Wikipedia, partisan reports, or secondary journalism is cited, it is as a pointer to primary sources or as evidence of public discourse, not as an authoritative claim. Readers are encouraged to verify claims against the primary sources linked. IMF data has been updated to the April 2026 Fiscal Monitor where relevant.

Appendix A

Formal Model of Verification Asymmetry & VerifyPrice

A.1 Definition 1: VerifyPrice(W) and Verification Overhead

For each canonical workload W (for example, a constraint system PROOF_{20} , a matrix multiplication size MATMUL_{4096} , or a verified-inference task $\text{INFER_LM}_{70\text{B}_{256\text{TOK}}}$), we care about two things:

1. How expensive it is to verify a purported result on a reference verifier, and
2. How that verification cost compares to producing the result from scratch.

VerifyPrice(W): verification profile Fix a reference verifier and adversarial mix of inputs. From a sample of verification runs for workload W , we estimate:

- $p_{50,t}(W), p_{95,t}(W)$: median and 95th-percentile wall-clock verify time from “artifact received” to “accept/reject”;
- $p_{50,c}(W), p_{95,c}(W)$: corresponding median and 95th-percentile verify cost in a reference unit (e.g., USD of cloud compute);
- $\text{fail}(W)$: observed fraction of artifacts that fail verification (invalid, malformed, or timed out).

The VerifyPrice vector for W is then:

$$\text{VerifyPrice}(W) \equiv (p_{50,t}(W), p_{95,t}(W), p_{50,c}(W), p_{95,c}(W), \text{fail}(W))$$

Canonical form and its relations. This five-field tuple is the canonical definition of VerifyPrice, and the variant tuples elsewhere in the thesis are views of it, not competing definitions. The physical-component views — $(t_{50}, t_{95}, e, m, f)$ in [Section 20.6](#) and $(t_{50}, t_{95}, m, b, e, f)$ in [Section 24.3](#) — carry the measured physical quantities from which the cost percentiles are derived. The derivation is *per run, then percentiled*: each run’s cost is $c_{\text{run}} = (e_{\text{run}}/3.6 \times 10^6) p_{\text{energy}} + t_{\text{run}} p_{\text{opp}}$ with both reference prices pinned and republished on change ([Section 20.6](#)); $p_{50,c}$ and $p_{95,c}$ are then the percentiles of c_{run} across the window. The canonical vector publishes time percentiles, derived cost percentiles, and the failure rate; the physical-component views pub-

lish the inputs. A dashboard may show either, but both must reconcile to the same underlying receipts, and any SLO that names VerifyPrice (Section 20.7) reads on the canonical five-field form.

ProduceProfile(W): production profile Similarly, we can describe the prover’s cost to produce correct outputs for W . From a sample of production runs (or cost models) we estimate median and 95th-percentile produce time and cost:

$$\text{ProduceProfile}(W) \equiv (p_{50,t}^{\text{prod}}(W), p_{95,t}^{\text{prod}}(W), p_{50,c}^{\text{prod}}(W), p_{95,c}^{\text{prod}}(W))$$

The denominator is an attack surface, and gets its own contract. The verification *numerator* of $r(W)$ is produced under the multi-operator, signed-batch, adversarial-corpus regime of Section 20.6; the production *denominator* has no equivalent protection as written, and it needs one, because a prover who inflates reported production cost drives $r(W)$ down and the network looks more verification-asymmetric than it is. The measurement contract for ProduceProfile(W) is therefore the mirror of the numerator’s:

- **Reference production cost model per workload.** Production cost is estimated against a published reference model — pinned reference hardware, pinned kernel and stack versions, pinned energy price — exactly as the FP32 acceptance bound pins its reference kernel (Section 20.3). Self-reported production cost may be published but does not enter $r(W)$.
- **Independent estimation.** At least two independent estimators (distinct from the prover being measured) produce the reference-model estimate; divergence above 5% triggers the same investigation protocol as verification measurements.
- **Published derivation.** The mapping from reference hardware, energy price, and runtime to estimated production cost is public and reproducible, so any observer can re-derive the denominator from its inputs.
- **Asymmetry acknowledged.** Where independent estimation is impossible for a workload, $r(W)$ is published as *prover-favorable and unverified*, and cannot by itself support a passing reading on chain link 3 (Section 7.1).

The consequence for the SLO below is that a passing $r(W) \leq 0.3$ is a *joint* claim about both halves of the ratio: verification is cheap under the adversarial harness, *and* production cost is what the reference model says it is, not what the cheapest prover reports.

Verification overhead $r(W)$ We define separate time- and cost-based overhead ratios:

$$(v/p)_W^{\text{time}} = \frac{p_{50,t}^{\text{prod}}(W)}{p_{50,t}^{\text{prod}}(W)}, \quad (v/p)_W^{\text{cost}} = \frac{p_{50,c}^{\text{prod}}(W)}{p_{50,c}^{\text{prod}}(W)}$$

For convenience, we can summarize these as a single scalar:

$$r(W) \equiv \max \left((v/p)_W^{\text{time}}, (v/p)_W^{\text{cost}} \right)$$

A.1.1 Implementation Note: VerifyPrice Telemetry Schema (Sketch)

The exact encoding of VerifyPrice is implementation-dependent, but for operational dashboards and APIs it helps to standardize a minimal JSON-style schema per canonical workload. A skeletal example:

```
{
  "workload_id": "PROOF_2^20",
  "window_start": "2025-06-01T00:00:00Z",
  "window_end": "2025-06-07T00:00:00Z",

  "p50_verify_time_sec": 0.40,
  "p95_verify_time_sec": 0.90,

  "p50_cost_usd": 0.00030,
  "p95_cost_usd": 0.00055,

  "failure_rate": 0.0004,

  "r_time": 0.18,
  "r_cost": 0.21,
  "r_max": 0.21,

  "sample_size": 125000,
  "hardware_profiles": ["open_tee_v2", "gpu_cluster_Y_v1"]
}
```

This schema is not normative; it is illustrative. The key properties are:

- Each canonical workload (workload_id) has its own series of VerifyPrice snapshots over sliding windows.
- VerifyPrice is expressed both in time and cost percentiles, with the overhead ratio $r(W)$ reported explicitly.
- Enough context is included to interpret shifts (sample size, hardware profile mix, time window), but no user-identifying data is required.

Clients and observatories can compute VerifyPrice from raw receipts and emit them in this or a similar format. The important thing is that:

1. The mapping from receipts $\rightarrow$ VerifyPrice is public and reproducible, and

2. The resulting metrics are exposed in a way that any third party can verify and alert on.

A system that claims “cheap, public verification” but cannot produce machine-readable VerifyPrice series for its canonical workloads is, by construction, asking you to trust what it cannot or will not measure.

We say that W is **verification-asymmetric** if $r(W) \ll 1$: checking honest work is much cheaper than re-doing it from scratch, both in time and cost.

A network (or facility) satisfies the VerifyPrice SLO for workload W if, for a published target (e.g., $p_{95,t}(W) < 5s$, $p_{95,c}(W) < 10^{-3}$ USD), it maintains:

- $\text{fail}(W)$ below a small, explicit threshold, and
- $r(W) \leq 0.3$ (with a stretch goal $r(W) \leq 0.1$),
over adversarial mixes of inputs and artifacts.

A.2 Definition 2: Facility Energy Receipt (FER)

A Facility Energy Receipt (FER) is a signed, machine-readable summary of a site’s energy use and availability over a fixed time interval Δt . Each FER is identified by a unique hash `fer_id` and contains at least:

- **facility_id**: a stable identifier for the site or plant;
- **interval_start, interval_end**: timestamps delimiting the reporting window;
- **energy_in_kwh**: total electrical energy consumed at the site boundary over Δt ;
- **energy_it_kwh**: energy delivered to IT loads (provers, routers, storage);
- **heat_reused_kwh**: energy captured and reused (for ERE metrics);
- **pue, ere, wue**: standard efficiency and water-use indicators over Δt ;
- **outages**: a list of curtailment or outage events during Δt ;
- **carbon_intensity_kg_per_kwh**: average emissions factor for supplied energy;
- **signature**: one or more digital signatures attesting to the report’s correctness.

FERs are produced at regular cadence (e.g., every 5 minutes or hour). Other receipts (e.g., PIDL receipts for proofs) can reference a FER via `fer_id`, allowing allocators, auditors, and protocols to reconstruct the energy context in which a unit of verified work was performed.

A.3 Definition 3: PIDL Receipt (Proof Interface Definition Language)

A PIDL receipt is the minimal, canonical record of a verified interaction in the Create/Compute $\rightarrow$ Prove $\rightarrow$ Settle $\rightarrow$ Verify loop. It binds together the claim, the evidence, and the operational context into a single, serializable object.

At a minimum, a PIDL receipt contains:

- **receipt_id**: a unique identifier (e.g., a hash of the serialized payload);
- **claim_id**: an identifier or hash of the claim being attested;

- **workload_id**: a canonical tag for the workload type and parameters (e.g., PROOF_2²⁰, MAT-MUL_4096);
- **proof_ref**: a hash or pointer to the proof or transcript artifact;
- **sla_tier**: the service class under which the work was promised (e.g., Bronze/Silver/Gold);
- **verify_time_sec, verify_cost_unit**: measured time and cost for verification;
- **result**: an outcome flag (e.g., accept, reject, timeout);
- **optional context**: hardware_profile, fer_id, settlement_txid, policy/jurisdiction tags;
- **signatures**: digital signatures from the prover, verifier, and (if applicable) broker or router.

PIDL receipts are designed to be portable and composable: they can be embedded in blocks, stored off-chain with only their hashes committed on-chain, included in higher-level proofs, or presented to auditors and counterparties.

A.4 Definition 4: VerifyReach(N, R)

Reachability and path diversity for independent verifiers in network N and region R .

For a given network configuration N , client region R , and a target time budget T (e.g., “within 30 seconds”), VerifyReach summarizes how reliably an honest client can reach at least one honest verifier within T seconds over multiple transport classes.

The VerifyReach KPI for network N in region R is the vector:

$$\text{VerifyReach}(N, R) \equiv (p_{50, \text{conn}}(N, R), p_{95, \text{conn}}(N, R), \text{succ}_1(N, R), \text{succ}_2(N, R), \text{fail}_{\text{mix}}(N, R))$$

Where:

- $p_{50, \text{conn}}, p_{95, \text{conn}}$: median and 95th-percentile time-to-first-connection;
- succ_1 : 1-path reachability rate;
- succ_2 : multi-path reachability rate (resilience to single-path blocking);
- fail_{mix} : distribution over connection failure causes.

Example SLOs:

- $\text{succ}_1(N, R) \geq 0.98$ and $\text{succ}_2(N, R) \geq 0.90$ for major regions;
- $p_{95, \text{conn}}(N, R)$ under a few seconds or tens of seconds (depending on transport mix).

A.5 Definition 5: VerifySettle(C)

Time, reliability, and privacy quality of non-custodial settlement on a corridor C .

For a given settlement corridor C (e.g., BTC $\leftrightarrow$ XMR atomic swaps), VerifySettle(C) summarizes whether non-custodial settlement actually clears under its advertised SLOs.

The VerifySettle KPI is:

$$\text{VerifySettle}(C) \equiv (p_{50,t}(C), p_{95,t}(C), s(C), r_{\text{safe}}(C), f_{\text{mix}}(C), a(C))$$

Where:

- $s(C)$ (success): fraction of settlements that complete as intended within the SLA window;
- $r_{\text{safe}}(C)$ (refund-safe): fraction of failed attempts where all funds return safely;
- $f_{\text{mix}}(C)$ (failure mix): distribution of failure causes;
- $a(C)$ (anonset): summary of anonymity-set size and churn.

A corridor satisfies its VerifySettle SLO if:

- $\text{success}(C) \geq \text{target}$ (e.g., 0.95), and
- $\text{refund_safe}(C) = 1.0$ over adversarial mixes,

with no hidden custodian or discretionary freeze paths.

A.6 Definition 6: Repression Wedge

The repression wedge is the realized real return on rate-capped instruments when nominal yields are held below inflation:

$$r_{\text{real},t} \equiv \frac{1 + i_{\text{cap},t}}{1 + \pi_t} - 1 \approx i_{\text{cap},t} - \pi_t$$

A negative $r_{\text{real},t}$ indicates that holders of capped instruments are being taxed in real terms via financial repression rather than explicit default.

A.7 Definition 7: Liquidation Effect

The liquidation effect measures the effective “tax” on bondholders from negative real rates (after Reinhart–Sbrancia):

$$\text{Liquidation}_t \approx \max(0, -r_{\text{real},t}) \times \frac{\text{Domestic Gov't Debt}_t}{\text{GDP}_t}$$

Advanced economies saw negative real rates $\sim 1/2$ the time, with average savings $\sim 1\text{--}5\%$ of GDP/yr during 1945–1980.

A.8 Definition 8: Physical VerifyPrice

Physical VerifyPrice extends VerifyPrice(W) (Definition 1) from proof/compute claims to the physical infrastructure claims (FCRs, Appendix A.2) that back them. For a facility profile F , it is the vector

$$\text{PhysicalVerifyPrice}(F) \equiv (p_{50,t}^{\text{FCR}}(F), p_{95,t}^{\text{FCR}}(F), p_{50,c}^{\text{FCR}}(F), p_{95,c}^{\text{FCR}}(F), \\ \text{ChallengeSuccess}(F), \text{OpenDiscrepancies}(F))$$

where $p_{\cdot,t}^{\text{FCR}}$ and $p_{\cdot,c}^{\text{FCR}}$ are time and cost percentiles to independently verify F 's FCR claims (energy profile, outage/curtailment history, hardware-profile integrity, jurisdictional routing), $\text{ChallengeSuccess}(F)$ is the fraction of third-party audit challenges that resolve without dispute, and $\text{OpenDiscrepancies}(F)$ is the count of unresolved claim mismatches. A Work Credit is high-assurance only if $\text{VerifyPrice}(W)$, $\text{PhysicalVerifyPrice}(F)$, and the relevant DVC bound clear their respective SLOs; none of those conditions makes the credit monetary. See [Sections 15.9](#) and [20.8](#).

A.9 Definition 9: VerifyFlow

VerifyFlow is the market-realization counterpart to VerifyPrice , VerifyReach , and VerifySettle . Where those three measure whether the stack supplies triad capacity under adversarial conditions, VerifyFlow measures the external financial representation and price-transmission state of the native monetary object. It is a pair of vectors, one per wrapper and one per asset:

$$\text{VerifyFlow}_i \equiv (L_i, A_i, \Delta A_i, \varepsilon_i, \kappa_i, C_i, H_i, D_i)$$

$$\text{VerifyFlow}(X) \equiv (Q^{RC}, Q^{RD}, \text{MPR}, \text{CCR}, \text{SER}, \text{WNG}, \text{RCR}, \text{NUS})$$

Component definitions appear in [Section 24.8](#); the estimating equations and their measurement contract appear in [Appendix H](#). VerifyFlow is explicitly *not* a protocol SLO: no VerifyFlow value falsifies the monetary thesis by itself. Its function is to determine whether observed price carries any information about native adoption ([Corollary 11.1](#)).

Appendix B

Practical KPIs & Telemetry Templates

This appendix provides a practical checklist of key performance indicators and telemetry metrics organized by operational domain. These templates support the VerifyPrice, VerifyReach, and VerifySettle observability regime described in the main text.

Prover markets: Queue depth, clearing price per proof type, cancellation rates, p50/p95 VerifyPrice.

Compute consensus pilots: MatMul throughput (GFLOPs/s), verifier cost ratio, invalid-work rejection rate under adversarial load.

Privacy settlement: Atomic swap success %, time-to-finality, stuck-flow causes, wallet UX friction metrics.

Decentralization & capture risk:

- Nakamoto coeff. (by stake/hash/prover share); top-N concentration; Gini.
- Entry latency: time-to-first-proof/mined-block for a new node.
- Geographic/ASN diversity; % over Tor/VPN; relay/MEV-builder diversity.

Liveness & safety:

- Reorg/orphan rate; effective finality time (p50/p95); incident MTTR/MTBF.
- SLA attainment (% proofs/settlements meeting latency/uptime targets).
- DoS/churn resilience: success % under adversarial load mixes.

Verify/produce economics:

- Verifier cost ratio (verify/produce) by workload; energy per proof/inference.
- Hardware mix & utilization (GPU/CPU/ASIC) vs. throughput; cost/FLOP (\$/GFLOP-verified).
- Failure taxonomy: prover faults, invalid proofs, circuit timeouts.

Privacy rails health:

- Anonymity-set size & churn (shielded pool/epoch); linkability regressions found.
- Swap slippage & fees (p50/p95); cross-venue route diversity; K-fail root causes.

- UX friction: successful first-run rate, steps/clicks to complete, abandonment rate.

Bridges & cross-chain settlement:

- Atomicity violations (0-conf leaks, partial completes); retry rate; liquidity depth per route.
- Watchtower/guardian coverage; light-client verification share vs. multisig.

Provenance & policy demand:

- C2PA attach rate, strip rate by platform; “verified views” share.
- Proof volume indexed to policy milestones (EU AI Act dates); % budgeted vs. ad-hoc spend.

Security & abuse:

- Sybil detection hit rate; staking/miner collusion alerts; MEV/censorship events.
- Key/attestation hygiene: validator/prover key rotations, slashing/penalties.

Token & issuance telemetry:

- Realized issuance vs. schedule; fee/burn coverage of security budget; % rewards tied to useful work.
- Velocity vs. locked/collateralized supply; SoV signals (holder age bands, exchange outflows).

Customer success:

- Enterprise proof spend: \$/month, SLA breaches, churn; Net Proof Retention (NPR).
- Time-to-integration (SDK→prod); support tickets per 1k proofs/settlements.

Hardware honesty:

- Share of receipts tied to open or sampled hardware profiles (by workload and volume).
- Lot-sampling coverage: percentage of lots and devices subject to destructive audits (SEM/optical imaging, side-channel tests), and maximum time since the last audit per profile.
- Incidents and deprecations: count and severity of profile-level compromises, time-to-deprecation, and migration progress away from affected profiles.

Appendix C

The SDK (Proofs-as-a-Library)

At the top of the stack, the application layer, you work with a small vocabulary:

- “Prove this computation.”
- “Prove this provenance.”
- “Settle this payment under these privacy and finality constraints.”
- “Anchor this result to hardware with this security profile.”

The SDK turns those sentences into types. A claim is a first-class object: `ComputationClaim`, `ProvenanceClaim`, `SettlementClaim`, optionally decorated with requirements. The compiler and runtime then map those claims to whatever combination of proving backends, settlement rails, and verifiable machines currently clear the SLA at the best `VerifyPrice`.

Core types:

```
class Policy:
    max_verify_time_sec: float
    max_verify_cost_usd: float
    privacy_level: str          # e.g. "encrypted_io", "public"
    finality_target: str        # e.g. "1m", "30m", "2h"
    allowed_hardware_profiles: list # e.g. ["open_tee_v2", "pure_zk_only"]

class ComputationClaim:
    fn: Callable
    inputs: Any
    policy: Policy

class ProvenanceClaim:
    media_bytes: bytes
    device_profile: str
    policy: Policy

class SettlementClaim:
    payments: list # [(recipient, amount, asset), ...]
    policy: Policy

class Receipt:
```

```

claim_id: str
proof_bytes: bytes
settlement_txid: str | None
metadata: dict # verify_time, verify_cost, hw_profile, etc.

```

SDK surface:

```

sdk = ProofSDK(networks=[...]) # PoW chains, zk rollups, open-TEE clusters

claim = sdk.make_claim(...)
receipt = sdk.prove_and_settle(claim, pay_with=user_wallet)
result = sdk.verify(receipt) # -> {accept: bool, metadata: ...}

```

Everything else—choice of proving backend, choice of settlement rail, choice of hardware profile—is a routing decision made under the hood, constrained only by the policy you declared.

Example 1: Proofed medical inference on verifiable machines

```

policy = Policy(
    max_verify_time_sec = 1.0,
    max_verify_cost_usd = 0.001,
    privacy_level = "encrypted_io",
    finality_target = "30s",
    allowed_hardware_profiles = ["open_tee_v2", "pure_zk_only"]
)

@proofed(policy=policy)
def diagnose(image: EncryptedImage) -> Diagnosis:
    return model.predict(image)

# In request handler:
claim = sdk.make_computation_claim(fn=diagnose, inputs=enc_image, policy=policy)
receipt = sdk.prove_and_settle(claim, pay_with=clinic_wallet)
result = sdk.verify(receipt)

```

Example 2: Camera provenance anchored in open hardware

```

# On capture (running on "open_camera_v1" device):
raw_bytes = camera.capture()

policy = Policy(
    max_verify_time_sec = 0.5,
    max_verify_cost_usd = 0.0005,

```

```
    privacy_level = "hide_location_and_identity",
    finality_target = "5m",
    allowed_hardware_profiles = ["open_camera_v1"]
)

prov_claim = sdk.make_provenance_claim(
    media_bytes=raw_bytes,
    device_profile="open_camera_v1",
    policy=policy
)

receipt = sdk.prove_and_settle(prov_claim, pay_with=newsroom_wallet)
bundle = MediaBundle(media=raw_bytes, provenance_receipt=receipt.export())
```

Example 3: Private payroll over neutral rails

```
policy = Policy(
    max_verify_time_sec = 3.0,
    max_verify_cost_usd = 0.002,
    privacy_level = "shielded_settlement_with_auditable_receipts",
    finality_target = "2h",
    allowed_hardware_profiles = ["pure_zk_only"]
)

payroll_batch = [
    Payment(recipient=alice_addr, amount=1000, asset="USDt"),
    Payment(recipient=bob_addr, amount=1200, asset="USDt"),
]

settle_claim = sdk.make_settlement_claim(payments=payroll_batch, policy=policy)
receipt = sdk.prove_and_settle(settle_claim, pay_with=treasury_wallet)
```

Appendix D

Energy & Plant Architecture

Plant architecture for verifiable work Objective: deliver verifiable, dispatchable, and composable power to proving and inference clusters so that proof units and verified FLOPs remain cheap to check and neutral to route.

Electrical topology (MV $\rightarrow$ LV):

- Medium-voltage interconnect with dual utility feeds where available; 2N or N+1 step-down to rack power.
- UPS/BESS for ride-through and proof-preserving shutdowns.
- Power-quality SLOs (voltage/frequency/THD) bound worst-case verifier latency inflation.

Thermal topology:

- Air to liquid (direct-to-chip/immersion) as default for high-density provers.
- Heat-reuse loops to district heating, greenhouses, or absorption chillers to improve ERE.
- Publish PUE and ERE: if we monetize heat, $ERE < 1.0$ is both possible and economically material.

Network & provenance:

- Redundant backhaul with clock discipline (PTP/Stratum) so receipts carry tight timing jitter bounds.
- Facility attestation path: meters and controllers sign energy/thermal telemetry.

Control plane:

- Workload-aware orchestration couples grid signals $\rightarrow$ job scheduler.
- Interruptible Bronze proofs pause first; Gold proofs ride through with BESS and on-site firming.

Energy procurement & grid posture:

- Portfolio, not a point bet: long-dated PPAs, nodal index exposure, behind-the-meter firming, and BESS.

- Co-location with stranded or curtailed energy (hydro spill, wind/solar curtailment, flare-gas).
- Demand response as a feature: treat prover/inference farms as dispatchable loads.
- Carbon & provenance: bind certificates into FER format so “kgCO₂e per proof” is auditable.

SLA tiers that map power quality to receipts:

- **Bronze (interruptible):** Shed priority 1. Eligible for demand-response; refunds/penalties via SLA escrow if p95 VerifyPrice breached. Typical use: PoL audits, non-urgent prover workloads.
- **Silver (curtailable with notice):** Shed priority 2 with N+1 firming; curtailment windows posted ahead of time. Typical use: batch MatMul-PoUW, rollup proving backlogs.
- **Gold (firm):** 2N electrical path or N+1 + BESS ride-through; zero curtailment except force majeure. Typical use: settlement-critical proofs, receipt ledger finalization.

Appendix E

Hardware Profiles

Throughout this thesis we treat “hardware profiles” as first-class objects. They are the way we talk about machines in the same language we use for proofs and receipts.

E.1 Concept

A hardware profile is a public description of a class of devices that are interchangeable for security and performance purposes. It is not a single physical unit; it is the specification for what counts as a valid unit of that type.

Informally: a profile says, “devices that meet this description behave like this, within these error bars, under this sampling regime.”

We use hardware profiles to:

- express trust and threat-model assumptions at the protocol and application layer;
- bind proofs and receipts to the types of machines that produced them; and
- make the honesty (and failure rates) of machines measurable over time.

E.2 Naming

Profiles are identified by short, stable strings. Three broad kinds:

1. **Open-hardware profiles** — devices whose RTL/microarchitecture and sampling regime are publicly documented.
 - `open_tee_v2` — second-generation open TEE design with published RTL, side-channel budget, and lot-sampling plan.
 - `open_signer_v1` — simple key-storage/signing chip with open RTL and sampling.
 - `open_camera_v1` — camera/capture pipeline with documented sensor path, secure element, and sampling method.
2. **Logical or pure-cryptography profiles** — paths that do not rely on special hardware.
 - `pure_zk_only` — verification/proving path relying only on general-purpose CPUs/GPUs and cryptographic assumptions.
 - `software_only_v1` — profile used for testing or where hardware attestations are unavail-

able.

3. **Vendor or mixed profiles** — existing devices not fully open but characterized and sampled.
 - `vendor_tee_X_v1` — proprietary TEE profile with documented side-channel and attestation scheme.
 - `gpu_cluster_Y_v1` — particular GPU farm configuration with known performance and sampling behavior.

E.3 Hardware Profile Specification

Each profile has a specification that is part of the public “bill of materials” for the stack. At a minimum:

- **profile_id**: string (e.g., “open_tee_v2”)
- **device_class**: what the device is used for (e.g., “TEE”, “signer”, “accelerator”, “camera”)
- **trust_mode**: “open”, “mixed”, “opaque”, or “logical”
- **design_commit**: identifiers for design artifacts (hashes of RTL, GDS, firmware images)
- **side_channel_budget**: qualitative/quantitative bounds
- **sampling_plan**: how lots and devices are checked
- **attestation_scheme**: how devices prove identity and state
- **security_target**: high-level security goal

E.4 Hardware Profiles in Receipts and Policies

Profiles appear in three main places:

1. **Policies** — Developers express requirements using profiles:

```
allowed_hardware_profiles = ["open_tee_v2", "pure_zk_only"]
```

2. **Receipts (PIDL)** — Every PIDL receipt includes a `hardware_profile` field.
3. **Telemetry** — The observability layer aggregates metrics by profile.

E.5 Example Profiles

open_tee_v2:

- **device_class**: “TEE”
- **trust_mode**: “open”
- **design_commit**: hashes of RTL, layout, firmware
- **sampling_plan**: random lot sampling, SEM/optical inspection
- **intended use**: high-value proving, key storage, secure computation

open_camera_v1:

- device_class: “camera”
- trust_mode: “open”
- sampling_plan: device sampling and destructive testing per lot
- intended use: origin-anchored capture for media provenance

pure_zk_only:

- device_class: “logical”
- trust_mode: “logical” (no hardware assumptions)
- intended use: contexts where hardware trust is unacceptable

Appendix F

Communications Resilience Mechanisms

P2P transport set:

- v2 Encrypted P2P (BIP-324) as default for all full nodes and relays; expose QUIC/TLS fallbacks.
- Onion/I2P modes first-class in clients; auto-failover if clearnet handshakes exhibit DPI resets.
- Handshake camouflage (Noise/obfs-style) for relays in high-interdiction ASNs.

Receiver-private routing for payments:

- BOLT12 Offers (receiver-private, reusable invoices) in wallets and merchant stacks.
- Path-blinding where supported.
- Federated ecash and shielded pools as optional “last-mile” sinks/sources.

Settlement survivability:

- Adaptor-signature atomic swaps for BTC↔XMR/ZEC corridors with typed failure & refund flows.
- Wallet UX exposes “abort & refund” as a first-class path.

Topology & anti-eclipse hardening:

- Peer-set diversity targets (geo/ASN spread), randomized peer rotation, inbound slot quotas.
- Gossip-path multiplicity to reduce single-jurisdiction capture.
- Light-client modes that verify receipts over any reachable path (browser, mobile, enclave).

Appendix G

Glossary of Terms & Notation

Triad Privacy, Proofs, and Compute — three cryptographic service capacities that may support a monetary candidate. Whether they should share one asset is tested by the Triad Coherence Test, not assumed.

VerifyPrice(W) A public KPI vector measuring the cost and time to verify workload W .

VerifyReach(N, R) Metrics for network reachability under censorship conditions.

VerifySettle(C) Metrics for settlement success and safety on corridor C .

Value Capture Lemma The bridge from service use to asset accrual: demand for triad capacity reaches the native asset only if five conditions hold—required fee medium, supply reduction via burns, collateral lockup, issuance discipline, and non-bypassability. It establishes a cash-flow claim and floor, not monetary premium. See [Lemma 11.1](#).

$r(W)$ Verification overhead ratio — $v(W)/p(W)$, where v is verification cost and p is production cost.

Base Asset The protocol's bearer fee and settlement unit and the only instrument tested as a monetary candidate. Fees, burns, collateral, or governance use establish neither moneyness nor a durable holder anchor by themselves.

Work Credit A typed capacity or service claim bounded by stress-adjusted deliverable output. It may be workload-, location-, hardware-, SLA-, and time-specific; it is not presumed to be a store of value.

Bypass Risk The risk that users consume triad capacity (proofs, privacy, compute) through channels that do not require the native asset—e.g., fiat-denominated cloud services, stablecoin payment to operators. If bypass channels dominate, the SoV thesis fails.

Differential value (Δ) The marginal buyer's willingness to pay for the protocol's differential properties—cryptographic checkability, neutrality, credible non-discretion, non-custodial exit—over the best bypass channel. Bounds the sustainable fee: $f \leq \Delta - (c_p - c_b)$, where c_p and c_b are the protocol's and the bypass channel's all-in costs. Also written $\Delta(R)$ when its dependence on regime pressure is at issue. A scalar wedge in price units, not the difference operator. See [Section 11.10.3](#).

Fee incidence The question of who bears a protocol fee. With free entry and roughly constant returns, long-run supply is horizontal at marginal cost, so the fee is passed to buyers and operator economic profit is unchanged: a fee is a wedge on *turnover*, not a claim on residual profit.

With installed capacity fixed, part of the fee is instead absorbed as quasi-rent to deployed hardware, which affects entry. See [Section 11.10.1](#).

Capturable wedge (ω) $\omega = s - \gamma$, where s is the normalized bypass spread $(P^{\text{prot}} - P^{\text{byp}})/P^{\text{byp}}$ and γ is the verification cost share (proving plus redundancy) as a fraction of P^{byp} . The measurable proxy for Δ : what remains of the price the protocol commands after paying for the verification that justifies it. Normalized because the absolute spread shrinks mechanically with compute deflation. Read by Red Line 14, Condition A.

Realized take rate (τ) Native fees plus burns in USD, divided by settled protocol turnover in USD. Computable by any third party from chain data. Compared against ω rather than against a guessed threshold, which is what makes Red Line 14's first condition self-normalizing.

Regime-Contingent Convenience Yield The channel through which a triad asset could earn monetary premium: the *holder-side service flow* arising from the state-contingency of $\Delta(R)$ —the bearer's ability to transact, prove, hold, and exit when the substitutes for proofs, courts, and custodians have stopped working—plus credible non-discretion and bearer holdability. It accrues by virtue of holding rather than as a distribution, so there is no stream to discount, which is what places it outside a discounted-cash-flow valuation. Named by analogy to the convenience-yield *mechanism* of Kaldor (1939) and Working (1949), not by any claim that a triad asset is a storage commodity. Explicitly *not* the countercyclicality of the fee stream, which is a negative-beta discount-rate effect a DCF captures exactly; and explicitly *not* produced by burns, collateral, or fee share, which yield a cash-flow claim and a floor. Unsized. See [Section 11.10.8](#).

Collateral unit-elasticity The property that a collateral requirement specified as a share of capacity value locks a *value* rather than a quantity: required units are $N = \chi K / P$ for collateral ratio χ , capacity value K , and price P . Consequences: a level effect rather than a growth effect, a bound of χK set by an accounting identity, wrong-way risk in large drawdowns, and an effect on the liquidity state Λ_t rather than the monetary impulse θ_t . See [Section 11.10.9](#).

Utility-Token Trap The pattern where a network provides useful services but the native token fails to capture economic value because operators extract all surplus, users pay in alternative currencies, or governance inflates supply. The thesis's value-capture conditions are designed to prevent this outcome.

PIDL (PIDL) Proof Interface Definition Language — the minimal receipt schema for proofs and settlements.

PAL (PaL) Proofs-as-a-Library — SDK that compiles claims to proofs.

PRK (PRK) Privacy Rails Kit — executes non-custodial, refund-safe settlement over privacy corridors.

FER Facility Energy Receipt — signed summary of a site's energy use over a time interval.

FCR Facility Capacity Receipt — extends FER data with infrastructure resilience, grid, cooling, jurisdictional, and hardware-diversity metrics. See [Section 15.7](#).

Delivered Verified Capacity (DVC) Scenario-specific maximum surviving flow from energy and fuel through firm electricity, transformer and switchgear, cooling and water, hardware, network, workload, proof, verification, settlement, and usable service. Reports the active minimum cut, substitution latency, and common-cause dependencies. See [Section 15.9](#).

Physical VerifyPrice Used in two related senses, distinguished in [Section 20.8](#): (i) the constitutional SLO measuring verification cost in real resources on reference hardware, exogenous to token price; and (ii) the cost of auditing the physical infrastructure claims (FCR) behind a unit of verified work. Red Line 1 refers to sense (i); Red Line 10 to sense (ii).

Sovereign Optionality ($\mathcal{O}_s$) A capacity-weighted index of the number of independent physical pathways by which verified work can continue under disruption, net of coercion exposure. Computed from existing FCR fields; feeds risk haircuts and Red Line 13. See [Section 15.8](#).

Topological Scarcity Lemma A nominal stock does not secure a monetary service; the relevant quantity is stress-deliverable flow through the narrowest non-substitutable edge connecting input to holder utility. See [Section 7.2](#).

Pressure–Capacity Corridor The relation $\Pi_{\text{monetary}}(R) \propto \Delta(R) \times A_{\text{full-stack}}(R) \times H(R)$. Its likely shape is an inverted U because differential demand rises as substitutes weaken while stack availability and holder agency can fall under extreme pressure. See [Section 7.2](#).

Disruption-Adjusted VerifyPrice The probability-weighted cost of verification across physical disruption states, rather than the observed cost under benign conditions. Used to price whether a resilience premium is buying anything. See [Section 15.10.2](#).

Energy Interdiction The adversary class in which curtailment, rationing, tariff discrimination, interconnection denial, or long-lead equipment scarcity is used to raise the physical cost of verification without any prohibition on cryptography. See [Section 5.5](#).

PoUW Proof of Useful Work — consensus mechanism where block rewards are earned by producing verifiable receipts of useful compute.

SLO Service Level Objective — published targets for system performance and availability.

Bronze/Silver/Gold SLA tiers for proof and settlement services with different latency, redundancy, and interruptibility characteristics.

MatMul-PoUW Proof of Useful Work construction based on matrix multiplication verification.

ZK Money An analytical lens or legacy label for service instruments referencing Privacy + Proofs. Those claims are not presumed money; only the separate base asset is a conditional monetary candidate.

Attestation Money An analytical lens or legacy label for proof-capacity and provenance-attestation demand. Formerly “Truth Money” in earlier drafts; it does not promote proof claims into money.

AI Money An analytical lens or legacy label for service instruments referencing Compute + Proofs (verified FLOPs, inference capacity). Work Credits beneath the lens remain typed ser-

vice claims.

Layer 0 Verifiable Machines & Energy — open hardware and sampled supply chains as base reality.

Layer 1 Reachability — communications and transport resilience.

Layer 2 Distribution & Execution — software supply and runtime.

Layer 3 Identity & Claims — pseudonymous credentials without doxxing.

Layer 4 Truth & Work — proof systems, PoUW, VerifyPrice.

Layer 5 Value & Settlement — privacy rails and non-custodial flow.

Layer 6 Governance & Telemetry — keeping neutrality and resilience measurable.

Market Realization Plane The external financial machinery—exchanges, custody, ETFs, treasury vehicles, index products, derivatives, dealers, passive mandates, systematic and agentic allocation rules—through which claims on the native monetary object are represented and priced. Orthogonal to the stack, not a layer within it. See [Section 11.14](#).

Native instrument A protocol-internal asset, typed service claim, evidence artifact, staking or LP position, corridor claim, or capacity voucher. Touching the protocol does not make the instrument monetary.

External financial wrapper A conventional market product written on the asset (spot ETF, ETP, treasury company, custodial balance, future, option, swap, leveraged or inverse ETP). Holding it exercises nothing. Unqualified “wrapper” means this.

Value capture vs. price capture Value capture routes native demand through the monetary object via fees, burns, and collateral. Price capture is demand for exposure to the asset’s price, which can occur with no protocol use at all. Neither implies the other. See [Corollary 11.1](#).

VerifyFlow The fourth verification family: telemetry for the external financial representation and price-transmission state of the native asset. See [Appendix A.9, Section 24.8](#).

Flow elasticity (ϵ) The responsiveness of a wrapper’s shares outstanding to changes in its value per share. $\epsilon \approx 0$ means holders sit still; $\epsilon \approx -1$ means they trim to a constant dollar position. Measured, not structural. See [Appendix H.1](#).

Net mechanical gain (κ) $L(L-1) + \epsilon L^2$; the coefficient converting an underlying return into a wrapper’s total return-coupled exposure demand. See [Appendix H.3](#).

Wrapper Recycling Ratio (WRR) $-\epsilon L/(L-1)$. Equal to 1 when holder flows fully offset gross rebalancing; below 1 the wrapper amplifies; above 1 it is countercyclical. See [Appendix H.5](#).

Return-coupled vs. return-decoupled flow Return-coupled flow (Q^{RC}) responds to the return itself and shapes volatility and persistence. Return-decoupled flow (Q^{RD}) arrives largely independent of the day’s return and shapes destination and level. See [Appendix H.7](#).

Wrapper–Native Growth Gap (WNG) $\Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$, computed on unit-denominated quantities over a trailing four-quarter window (asset-price revaluation stripped from both sides), with the USD-denominated series published alongside. Persistently positive means financial exposure is outgrowing monetary use.

Service-Good Realization The pattern in which triad capacity is consumed at scale while the asset shows no insurance signature across repeated regime-pressure episodes and the monetary bid accrues elsewhere. Prescribes reclassification to *verified-capacity service asset*, not retirement. Filed as Market Realization Warning 13 because its corroborating clause reads on price; Red Line 15 separately covers the native collateral–capacity spiral. See [Section 28.7](#).

Exhaustion gradient The approach path to the participation line: a household remains formally inside the system (banked, employed, housed, consuming) while its capacity to accumulate declines, so aggregates register nothing until a delinquency or unemployment threshold is crossed. Dollar-weighted measures lag person-weighted depletion by construction. See [Section 3.6](#).

Household transmission The layer converting aggregate growth into household income, security, shelter, and savings. Its characteristic failure is depletion beneath resilient aggregates: consumption sustained by drawdown rather than income. Formally, operating surplus $O_i = Y_i + I_i + T_i - E_i - DS_i < 0$ financed by depleting resilience stock R_i . See [Section 30.9](#).

Distributed resilience The requirement that system resilience be assessed jointly across the sovereign layer (R_{state} : energy, industry, defense, administration) and the household layer ($R_{\text{household}}$: shelter, savings, family formation, shock absorption), because a state can remain stable while financing continuity from household buffers — $R_{\text{state}} \neq R_{\text{household}}$. Fragility transferred downward returns through politics and defaults. See [Section 30.9](#).

Representation discipline The publication rule that every aggregate over households or holders discloses its weighting rule, because dollar-weighted and person-weighted summaries of the same observations can move in opposite directions: $\bar{x}_{\text{dollar}} - \bar{x}_{\text{person}} = \text{Cov}(C_i, x_i)/\bar{C}$. A verified aggregate can be accurate and describe almost no participant; the weights are the explanation, and they are published. See [Section 24.8.8](#).

Capability vs. formal access Formal access is a valid balance and the legal right to transact; capability is command over the necessities of an independent life. A ledger can supply the first while the second fails, because making claims harder to alter does not broaden who can acquire them. The ninth SoV requirement is enforced on both layers. See [Item 9, Section 31.3.5](#).

Capital Survival Ratio (CSR) Whether successive investor cohorts in a wrapper preserved capital, as distinct from whether the product preserved AUM. See [Section 24.8.4](#).

Policy Concentration Ratio The share of net asset demand governed by the largest common rebalancing templates, treasury algorithms, or agent policies. Ownership can be decentralized while behavior is not.

Compositional adversary A failure mode with no malicious actor: locally rational agents whose incompatible operating rules combine into concentration, procyclicality, recursive leverage, and uninformative prices. See [Section 5.2](#).

Collateral Loop The reflexive circuit in which asset values support tax receipts, receipts support fiscal capacity, fiscal capacity and central-bank support sustain Treasury market function-

ing, and Treasury yields sustain asset values. See [Section 3.2](#).

Collateralized Sovereign Stack A description of the incumbent macro order, in which a state's effective fiscal capacity depends on the market value of the asset complexes it regulates. A diagnosis, not an architecture; distinct from the seven layers this thesis specifies. See [Section 3.2](#).

Sovereign Credit Safety vs. Collateral Stability Credit safety is the probability of nominal repayment; collateral stability is the reliability of the instrument's price at the horizon over which it is used as margin and reference. An obligation can have the first and lack the second. See [Section 3.1.2](#).

Automatic vs. Marginal Buyer The automatic buyer supplies demand by formula. The marginal buyer absorbs residual duration and therefore sets the price. Quantity of bids and quality of risk absorption are different objects. See [Section 3.1.2](#).

DV01 The dollar change in value produced by a one-basis-point change in yield. The relevant unit for whether a buyer is warehousing interest-rate risk. See [Section 3.1.2](#).

Duration of the Claim vs. Duration of the Project Duration of the claim is the interest-rate exposure of a promised cash flow, which a repression-resistant store of value must not be. Duration of the project is the years between committing present resources and receiving the output of a long-lived physical asset. See [Section 3.3](#).

Duration Warehouse An institution capable of converting present savings into long-lived productive capacity without being forced to liquidate, refinance, or reprice at the worst moment. A diagnosis of the incumbent order, not a layer of the cypherpunk stack. See [Section 3.3](#), [Section 31.5.1](#).

Buyer-Quality Checklist Persistence of funding, investment horizon, match to an enduring liability, capacity to bear losses without forced liquidation, and countercyclical willingness to add risk when prices fall. Published as fields, not as a multiplicative index. See [Section 24.8.7](#).

Native Monetary Buyer Map Telemetry separating just-in-time fee acquisition, operator sell-through and inventory, burns and net issuance, self-custodied reserve demand, wrapper demand, leverage, holding period, loss-bearing capacity, and countercyclical accumulation. See [Section 24.8.8](#).

Project Note Explicit duration-bearing infrastructure credit with stated maturity, covenants, default states, and loss priority. Proofs may audit it; no par, redemption, or emergency-support promise may transfer to the base asset. See [Section 31.5.2](#).

Sovereign Maturity Transformation Financing long-lived public assets and commitments with short-term public liabilities. Duration is transferred onto the rollover calendar, not extinguished. See [Section 3.2](#).

National-Champion Pathway The conversion of private firms into protected quasi-public infrastructure because the state's fiscal or security position depends on them. Also *enclosure by rescue*. See [Section 5.2](#).

Numeraire-Dependence The property that a measured price path is jointly determined by the

asset and by the unit it is quoted in, so the same series can support opposite conclusions about real command over resources. See [Section 11.14.2](#).

Open Technology vs. Open Sovereignty Open-source or open-weight software is sovereign only where users can independently obtain the power, hardware, data, communications, privacy, and settlement to run it. Openness at one layer is compatible with concentration beneath it. See [Section 30.4](#).

Strategic Tempo The conversion of physical resilience into control over *when* a disruption becomes binding, and thence into bargaining power. See [Section 30.1](#).

Appendix H

Formal Model of Market Realization, Wrapper Flows, and Price Capture

This appendix develops the flow machinery behind the Market Realization Plane (Section 11.14) and VerifyFlow (Section 24.8). It is deliberately quarantined here: the main body needs the distinction between value capture and price capture, not the derivations.

Attribution. The flow machinery below—flow elasticity, holder recycling, leveraged-rebalancing mechanics, and the return-coupled versus return-decoupled distinction—is adapted from Michael W. Green’s work on how product mechanics and holder behavior set prices, in particular *A Semi-Theory of Almost Everything* (Tier1 Alpha, July 2026). The market-impact form is the square-root “law” of execution cost from the market-microstructure literature, in which impact scales roughly with the square root of order size relative to available liquidity; it is not Gabaix and Koijen’s aggregate demand-elasticity multiplier, which does not appear in this model. Gabaix and Koijen (2021) is relied on only for the broader inelastic-markets premise—that flows move prices at all, because aggregate demand for financial assets is far less elastic than textbook treatments assume—which is why a flow-based telemetry layer is worth publishing. What is ours is the generalization from equity market structure to monetary telemetry: the claim that these quantities must be published as a distinct verification family (VerifyFlow) precisely so that price movement is not mistaken for monetary adoption. Any errors in that generalization are ours, not theirs.

A caution before the mathematics. These equations describe *transmission and magnitude*. They do not explain why a narrative ignites a behavioral shift in the first place. Flow elasticity in particular is measured, not explained. Treating any parameter below as a structural constant is a modeling error.

H.1 Flow Elasticity

For wrapper i , estimate the responsiveness of shares outstanding to changes in value per share:

$$\Delta \ln S_{i,t} = \alpha_i + \varepsilon_{i,t} \Delta \ln \text{NAV}_{i,t} + u_{i,t}$$

where $S_{i,t}$ is shares outstanding, $NAV_{i,t}$ is value per share, and $\varepsilon_{i,t}$ is **flow elasticity**. Interpretation: $\varepsilon \approx 0$ means holders largely sit still; $\varepsilon \approx -1$ means holders redeem roughly enough after gains to maintain a constant dollar position.

The regression has three parts and each is allocated to a different place downstream, so the decomposition must be stated rather than assumed. The ε -driven component is return-coupled and belongs in Q^{RC} below. The intercept α_i captures baseline creation or redemption that occurs regardless of the day's return — distribution growth, model-portfolio adoption, scheduled contributions — and therefore belongs in the return-decoupled term Q^{RD} , not in Q^{RC} . The residual $u_{i,t}$ stays unexplained and is published as its own series, never absorbed into either aggregate: Q^{RD} consumes α only. Assigning α_i to Q^{RC} would overstate momentum amplification; absorbing $u_{i,t}$ into either term would launder unexplained variation into a named flow. The published triple $(\hat{\varepsilon}, \hat{\alpha}, u)$ is the complete accounting; the aggregates consume two of the three.

Identification: the regression is simultaneous, and the bias is not benign. $\Delta \ln NAV_{i,t}$ is not exogenous to $\Delta \ln S_{i,t}$. Creations and redemptions force dealer hedging, hedging passes through the impact function of [Appendix H.8](#), and impact moves the underlying — which moves the NAV being used as the regressor. The OLS estimate of ε therefore suffers simultaneity bias, and its direction is the dangerous one: flows that amplify the move induce a correlation between share changes and NAV changes that OLS books as elasticity, so naively fitted $\hat{\varepsilon}$ overstates amplification precisely when amplification is present, and the fitted values then feed κ and WRR downstream. Three disciplines follow. First, prefer instruments for $\Delta \ln NAV$ that are exogenous to wrapper flow — underlying-index constituent news, macro release days, cross-venue price dislocation — over the raw regression. Second, where instruments are unavailable, publish the OLS estimate as an upper bound on amplification rather than as a point estimate. Third, never recycle a fitted $\hat{\varepsilon}$ back into the stress harness of [Appendix H.12](#) without reporting both the instrumented and OLS values; a harness whose elasticity was fitted on the same amplification it simulates will find it.

The second bias runs the other way, and it decides when the upper bound is safe to publish. For a continuously traded underlying, the wrapper's official NAV is computed on a stale or synchronized print, not on the price at which flows actually cleared. The regressor is then a noisy measurement of the true underlying return, and classical errors-in-variables attenuation pulls $\hat{\varepsilon}$ toward zero — understating amplification. The two biases are opposite in sign: simultaneity inflates $\hat{\varepsilon}$, staleness attenuates it, and which dominates is not a property of the estimator but of the underlying's liquidity and the wrapper's pricing convention. The upper-bound publishing rule above is therefore conditional, not absolute. A fourth discipline: publish alongside $\hat{\varepsilon}$ a staleness diagnostic — the divergence between the wrapper's pricing

print and the underlying's contemporaneous venue price — and treat the OLS value as an upper bound only where that divergence is small. Where staleness is material, the OLS estimate bounds nothing in either direction, and the reading must be published as *unidentified* rather than as a bound.

Because a daily-reset leveraged wrapper's NAV return is approximately $L_i r_t$ for underlying return r_t :

$$\frac{\Delta S_{i,t}}{S_{i,t}} \approx \varepsilon_{i,t} L_i r_t$$

H.2 Leveraged Rebalancing and Flow-Induced Exposure

A daily-reset fund must trade to restore its target exposure. The gross required rebalance is

$$Q_{i,t}^{\text{rebalance}} \approx L_i(L_i - 1)A_{i,t}r_t$$

where $A_{i,t}$ is fund assets. Note that $L(L - 1) > 0$ both for $L > 1$ and for negative leverage such as -2 or -3 : **long and inverse leveraged funds both mechanically chase the underlying move**. The inverse product is not a stabilizer.

Separately, creations and redemptions carry exposure equal to leverage times the dollar flow:

$$Q_{i,t}^{\text{flow}} \approx \varepsilon_{i,t} L_i^2 A_{i,t} r_t$$

Combining the two gives total return-coupled mechanical exposure:

$$Q_{i,t}^{RC} \approx [L_i(L_i - 1) + \varepsilon_{i,t} L_i^2] A_{i,t} r_t$$

H.3 Net Mechanical Gain

Define the **net mechanical gain coefficient**

$$\kappa_{i,t} = L_i(L_i - 1) + \varepsilon_{i,t} L_i^2$$

so that $Q_{i,t}^{RC} = \kappa_{i,t} A_{i,t} r_t$. The sign and magnitude of κ determine whether a wrapper amplifies or absorbs the move it sits on.

H.4 The Recycling Boundary

A wrapper fully offsets its own gross mechanical trade when $\kappa_i = 0$, which implies a critical elasticity

$$\varepsilon_i^* = -\frac{L_i - 1}{L_i}$$

Daily leverage L	Full-recycling elasticity ε^*
+2	−0.500
+3	−0.667
−2	−1.500
−3	−1.333

Table H.1: Illustrative recycling boundaries. Holder behavior stops being a vague qualitative concern and becomes an observable regime boundary.

This is the single most useful conversion in the appendix: it turns “holder behavior matters” into a number that can be measured, plotted, and breached. Empirically, a $3\times$ product whose elasticity sits near -0.68 is close to self-neutralizing; if that elasticity weakens toward -0.40 , the same product becomes a material momentum amplifier, and the underlying’s behavior can shift from mean-reverting toward higher volatility with greater trend persistence.

H.5 Wrapper Recycling Ratio

A more legible operator metric normalizes elasticity by its own boundary:

$$\text{WRR}_{i,t} = -\varepsilon_{i,t} \frac{L_i}{L_i - 1}$$

where $\text{WRR} = 1$ means holder flows fully offset gross rebalancing; $\text{WRR} < 1$ means the wrapper amplifies the underlying move; and $\text{WRR} > 1$ means holder flows more than offset it and become countercyclical. This belongs on the public Market Realization & Wrapper Board.

The ratio and the gain coefficient are one object, and stating the identity prevents misreadings of both:

$$\kappa_{i,t} = L_i(L_i - 1)(1 - \text{WRR}_{i,t})$$

so the unnormalized quantity and its normalized image always move together, and any dashboard showing one can be checked against the other.

The unlevered case is the dominant one, and WRR is undefined there. WRR divides by $L_i - 1$, so it does not exist for an unlevered wrapper ($L_i = 1$) — and the spot ETF, the instrument

most likely to become this asset's largest wrapper, is exactly that. The degeneracy is not a defect in the metric; it is the metric correctly reporting that an unlevered wrapper has no rebalancing obligation and therefore nothing to recycle. What remains meaningful at $L_i = 1$ is the elasticity itself: holder flows into a spot wrapper scale with $\varepsilon_{i,t}A_{i,t}r_t$ through the Q^{RC} term, with no rebalancing component to offset. The publication rule is therefore: **leveraged and inverse wrappers publish κ and WRR; unlevered wrappers publish ε and $\kappa = \varepsilon Ar$; and the board never carries a WRR column for $L = 1$** . A reader who wants a single cross-class reading should use $\kappa_{i,t}A_{i,t}r_t$ directly, which is defined everywhere.

H.6 Wrapper-Class Applicability

The machinery above was derived for fund-like wrappers with shares outstanding, NAV, and creation/redemption mechanics — the equity-ETF structure it was adapted from. Not every wrapper class that matters for a bearer asset has those observables, and applying an equation to a class that lacks its inputs does not produce a number; it produces a fabricated one. The publication contract is therefore class-conditional:

Two consequences. First, the perpetual-futures substitute matters most in practice: for a crypto-native asset, perps are plausibly the largest leveraged wrapper class, and the funding-rate-to-price response is the closest available analog to elasticity — publish it, but do not label it ε , because it is not estimated from the same regression and does not carry the same interpretation. Second, the convenience-yield instruments of [Section 24.8.6](#) already read on the derivative complex (funding, basis); the two families meet exactly here, and a reading published by one should reconcile with the other where both touch the same venue.

H.7 Return-Decoupled Allocation Flow

Not all flow responds to returns. A thematic or passive wrapper can receive creations largely independent of the current day's return, arriving on up days and down days alike. For underlying asset j :

$$Q_{j,t}^{RD} = \sum_k w_{jk,t} F_{k,t}^\alpha$$

where $F_{k,t}^\alpha$ is the return-decoupled component of net creation into wrapper k and $w_{jk,t}$ is that wrapper's exposure weight to j . For leveraged or derivative wrappers, $w_{jk,t}$ must be a **delta-equivalent** exposure weight rather than a portfolio accounting weight, or the allocation flow will be understated by roughly the leverage factor.

The input discipline is the decomposition rule of [Appendix H.1](#), applied here as a boundary: $F_{k,t}^\alpha = \alpha_k A_{k,t}$, the intercept component of the creation regression only. The residual $u_{k,t}$

Wrapper class	Principal observables	Applies	Degenerates / substitute
Spot ETF / fund	Shares outstanding, NAV, creations	ε , κ , Q^{RD} (α -component), CSR	WRR undefined ($L=1$); publish ε in its place
Leveraged / inverse ETF	Shares, NAV, daily reset, L	Full apparatus: ε , κ , WRR, recycling boundary	None — this is the native habitat
Perpetual futures	Open interest, funding rate, long/short account ratio	Funding-rate response to price (a <i>substitute</i> elasticity), open-interest flow into Q^{RD}	No shares or NAV: ε via $\Delta \ln S$ vs $\Delta \ln \text{NAV}$ is undefined; the leverage flows of Appendix H.2 do not arise (no reset obligation) — but position changes remain return-coupled and belong in a Q^{RC} analog built on open interest
Dated futures / forwards	Open interest by expiry, term structure	Roll flow into Q^{RD} ; basis into convenience-yield telemetry (Section 24.8.6)	Same as perps; basis, not share mechanics, is the informative quantity
OTC swaps / treasury-company exposure	Disclosures, periodic filings	Delta-adjusted exposure into Q^{RD} at disclosure frequency	Stale observables; publish with the lag and never interpolate silently
Custodial / exchange balance	On-chain custodian-tagged balances	CCR, NUS, Wrapper Dominance components	No flow elasticity at all — a stock, not a flow instrument

Table H.2: Wrapper-class applicability of the VerifyFlow apparatus. A class without the observables an equation consumes does not get a fabricated reading; it gets the substitute construction named in the last column, or nothing.

is published as its own series and enters no aggregate; the ε -driven component is *excluded* because it is already inside κ_i and therefore inside Q^{RC} . Feeding raw net creation into Q^{RD} would count the same dollars twice — once as elasticity-driven exposure and once as allocation — and would propagate the overstatement into MPR, the stress harness, and any red line that reads on them. The two equations of this appendix therefore partition demand by return-dependence: **Q^{RC} owns the ε -component; Q^{RD} owns α ; the residual u is published and consumed by nothing.** Where a wrapper's creation series cannot yet be decomposed, publish raw $F_{k,t}$ *flagged as undecomposed* so the aggregation layer can exclude it from Q^{RD} rather than silently double-count it.

The two flow types do different work. Return-coupled flow governs volatility, momentum, persistence, and reversal severity. Return-decoupled flow governs *destination and level*: which assets receive the structural bid, how concentrated it becomes, and which prices are accepted without valuation-sensitive selling. For a triad asset, return-decoupled flow could arrive from a spot ETF, a crypto index product, a corporate treasury mandate, an automated wealth-allocation model, a retirement default, an agentic treasury system, or a regulated “digital hard asset” basket — all potentially enormous for price and nearly irrelevant to native triad usage.

H.8 Market Impact and Liquidity

Flow becomes price only through a liquidity-dependent impact function. The standard empirical form has impact growing roughly with the square root of trade size relative to available liquidity:

$$I_{j,t} = Y_j \sigma_{j,t} \operatorname{sgn}(Q_{j,t}) \sqrt{\frac{|Q_{j,t}|}{V_{j,t}}}$$

where Y_j is an impact coefficient, $\sigma_{j,t}$ is volatility, $Q_{j,t}$ is exposure demand, and $V_{j,t}$ is executable liquidity over the relevant horizon. Splitting impact into permanent and temporary components with permanence share ρ_j :

$$\Delta \ln P_{j,t} = \theta_{j,t} + \rho_j I_{j,t} + (1 - \rho_j) I_{j,t}^{\text{temporary}} + \eta_{j,t}$$

Fitted permanence is case-specific. Import the structure; estimate Y_j , ρ_j , and $V_{j,t}$ asset by asset.

H.9 Volatility Drag on Leveraged Wrappers

For a daily-reset wrapper with leverage L on an underlying with drift μ and volatility σ , the continuous approximation to compound growth is

$$g_L \approx L\mu - \frac{L^2\sigma^2}{2} - (L-1)r_f - f$$

where r_f is the financing rate on the levered portion and f is the product's expense ratio. The first two terms are the standard continuous approximation; the last two are the carry the wrapper actually bears and are frequently omitted, which flatters leveraged products in exactly the high-rate environments where they are most costly to hold. Note also that this section's results, along with the recycling boundary and WRR of [Appendix H.4](#), apply to daily-reset products with $L \neq 0, 1$; an unlevered or non-resetting wrapper has no rebalancing obligation and no drag term.

The variance penalty grows with the *square* of leverage. At sufficiently high volatility a leveraged wrapper destroys capital even when the underlying has a positive average return. This is the formal basis for the classification in [Table 7.4](#): a daily-reset leveraged wrapper cannot inherit the underlying asset's store-of-value status, because its terminal value depends on the path and not merely the endpoints. It is a path-dependent trading instrument that happens to reference a monetary object.

H.10 Dealer Hedge Decomposition

Product-level exposure demand is not equal to spot orders. Decompose realized hedging by instrument — physical holdings, futures, total-return swaps, options — and track delta- and gamma-adjusted exposure per counterparty, average derivative maturity, collateral and margin requirements, and evidence of counterparty caps. A migration from swaps toward options is a candidate signal that dealer balance-sheet capacity is binding, with the attendant gamma feedback when dealers are short calls. Treat this as a hypothesis to test per asset, not a law.

H.11 Measurement Contract

VerifyFlow Measurement Contract

- Use daily or intraday shares outstanding, NAV, price, AUM, holdings, leverage, and derivative disclosures.
- Chain reverse splits and corporate actions before estimating anything.
- Estimate elasticity over multiple rolling windows, never a single fixed window.
- Use structural-break tests rather than assuming fixed holder behavior.
- Address the simultaneity of elasticity estimation explicitly: instrument $\Delta \ln \text{NAV}$ or publish the OLS value as an upper bound on amplification, and only where the staleness diagnostic of [Appendix H.1](#) shows wrapper pricing tracks the contemporaneous underlying ([Appendix H.1](#)).
- Report exposure demand separately from realized spot orders.
- Delta-adjust swaps and options.
- Publish top- N counterparty and constituent concentration.
- Separate return-coupled and return-decoupled flows.
- Compare market-exposure growth against native protocol-use growth (WNG).
- Publish confidence intervals and model residuals.
- Report failed hypotheses rather than silently dropping them.

That last clause is not decoration. A flow model that only publishes the specifications that worked is indistinguishable from a narrative, and this thesis has no standing to demand receipts from protocols while exempting its own econometrics.

H.12 Reference Stress-Harness Sketch

For each scenario in [Section 24.12.1](#): enumerate wrappers and their $(L_i, A_i, \varepsilon_i)$; compute κ_i and WRR_i ; simulate a return path; accumulate Q^{RC} and Q^{RD} ; map aggregate Q through the impact function given V ; propagate the resulting return back into the next period's rebalance and creation terms; and record MPR, WNG, and native-use series alongside price. The output of interest is never the simulated price. It is whether the native series moved at all.

H.13 Boundaries and Caveats

- This machinery is evidence about *market transmission*. It is not evidence that Privacy, Proofs, or Compute will earn monetary premium.
- Parameters fitted in one market (equities, semiconductors) do not transfer. Re-estimate ev-

everything.

- Unlevered wrappers are not adversaries by construction. The failure mode is the inability to distinguish accessibility, custody concentration, native use, synthetic leverage, and price-insensitive demand.
- Price is never a protocol control target. The purpose of this appendix is to explain price formation, not to defend a price.
- Market realization is orthogonal to the stack. Nothing here becomes a Layer 7.

H.14 Buyer Quality and Price Provenance

The vectors above measure *product* mechanics. They do not measure *buyer* quality. Green's later work on the Treasury long end makes the gap precise: a recurring dollar inflow can absorb less duration as prices fall, because market-value weights shrink, and the residual then clears with whoever is left. The same structure can appear in a monetary wrapper whose creations continue while the holders behind those creations are default-enrolled, unlevered, and duration-unaware—or leveraged, short-horizon, and one mark from a forced sale.

VerifyFlow therefore adds a buyer-quality checklist, not a new elasticity (Section 24.8.7). For any material source of demand, record persistence of funding, investment horizon, match to an enduring liability, capacity to bear losses without forced liquidation, and whether incremental risk absorption rises or falls after a drawdown. Do not publish a multiplicative “anchor quality” index. The dimensions are the contribution; collapsing them into a scalar would recreate the scoreboard problem this appendix exists to prevent.

Approximate new DV01 absorption by an index receiving inflow F_t is

$$A_t \approx F_t \sum_j w_{j,t} D_{j,t} 10^{-4}$$

where $w_{j,t}$ is the market-value weight of instrument j and $D_{j,t}$ its modified duration. As $w_{j,t}$ falls with price, A_t can decline even as F_t is unchanged. That identity is why “inflows remain healthy” is not evidence that the residual is shrinking.

The measurement contract of this appendix applies here without modification: publish the classification of each buyer, the fields that support it, and the cases in which the classification failed. A flow model that reports only that demand arrived, and not what kind of demand it was, has not instrumented the plane.

Appendix I

Scenario Analysis: The Collateral Loop Under Stress

Epistemic Status of This Appendix

Nothing in this appendix is load-bearing. It exists to hold material that is analytically useful and evidentially weaker than the standard the main text holds itself to, so that the two are never confused.

The scenarios below are traced causal chains, not forecasts. They are offered because a mechanism worth defending against does not need to be probable, and because the discipline of writing a scenario down is what makes it possible to say later that it did not happen. [Appendix I.4](#) lists the claims from the same source material that did *not* clear the bar, and why.

No red line in [Section 28.6](#) and no condition in [Value Capture Lemma](#) depends on anything in this appendix.

I.1 Correlated Loss: A Recession With Rising Long Yields

The developed-market crisis template assumes a particular sequence: equities fall, capital flees to safety, Treasury prices rise, long yields fall. The sovereign's borrowing costs decline exactly when its receipts do, which is what makes the template survivable.

[Section 3.1.2](#) gives a reason that sequence can invert. If a meaningful share of marginal demand for the reference safe asset comes from leveraged, repo-financed intermediaries, then a volatility shock does not merely fail to attract those buyers—it converts them into sellers, because their constraint is margin rather than conviction:

- equity volatility rises
 - leveraged funds de-gross
 - arbitrage positions unwind
 - cash Treasuries are sold
 - long yields rise

Running concurrently, and on a slower clock:

- recession
 - receipts fall + automatic spending rises
 - issuance increases
 - term premium rises

Green's buyer-quality argument supplies a third concurrent mechanism. Automatic retirement and index flows can continue in dollars while absorbing less DV01 as long-bond prices fall, and target-date rebalancing—a relative-performance rule—contributes little when equities and bonds fall by similar amounts:

- stocks fall + bonds fall
 - little target-date rebalancing
 - + less automatic DV01 absorption
 - greater burden on the residual buyer

The plausible path is therefore two-phase: an initial flight to safety pushes long yields down, and the balance-sheet and fiscal effects then reverse it. The scenario is not that this must happen. It is that the conditions under which it *can* happen are now assembled, and that a portfolio or protocol treasury which assumes the standard template is unhedged against the inversion.

What this would mean for the thesis. Section 3.1 describes a playbook running debt to repression to flight-to-neutrality. That chain is unaffected in direction. What changes is the transition: the flight stage may be preceded by an interval in which conventionally safe and conventionally risky assets fall together, which is precisely the interval in which an unlevered holder of neutral collateral is advantaged and a levered one is destroyed. This is an argument for the leverage discipline in Section 24.12.1, not for any particular asset.

I.2 The AI–Fiscal Collateral Loop

Section 3.2 argues that asset values are load-bearing for fiscal capacity. The sharpest form of that argument concerns a single sector, because concentration is what turns a sectoral repricing into a sovereign one.

The chain, traced:

1. Competitive pressure or higher financing costs slow the growth of AI-related valuations.
2. Slower appreciation impairs collateral values and refinancing, before any outright price decline. The failure point is the second derivative: a financing structure built on continuously rising collateral can break when appreciation merely decelerates.
3. Capital expenditure and equity prices fall.
4. Capital-gains and stock-compensation receipts fall with them.
5. Consumption weakens through the wealth effect; recession follows.
6. The federal deficit widens.
7. Treasury issuance increases.
8. Long-term yields rise.
9. Higher yields further impair AI financing, closing the loop.

Magnitude discipline. The originating commentary attaches large numbers to step 5 in particular. Those numbers are not used here and are listed in [Appendix I.4](#). The architecture is the contribution, and it does not require them: the loop is a loop at any magnitude, and the question a reader should carry away is whether the sector’s share of index concentration, household portfolios, stock-based compensation, and construction activity is large enough for step 4 to be fiscally visible. That is a measurable question, and [Appendix I.5](#) says where to look.

Why a monetary thesis cares. Not because it forecasts a bust. Because the loop is the mechanism that would convert a technology repricing into pressure for the state interventions described in [Appendix I.3](#)—and those interventions, not the repricing, are what threaten an open stack.

I.3 The National-Champion Conversion Pathway

Section 5.2 introduces enclosure by rescue as an adversary class. This section sketches what the instruments would plausibly be, because an adversary described only in the abstract cannot be monitored.

A rescue of a fiscally load-bearing technology sector would not resemble a 2008 bank recapitalization. Central banks can support Treasury market functioning, dealer balance sheets,

repo, bank funding, credit conditions, and general discount rates. They cannot underwrite a specific firm's margins, a private valuation, a proprietary moat, or the commercial viability of a particular facility. Support for those has to come from the fiscal and regulatory side, and it arrives as a portfolio. The scenario is not that any of this has been decided or assembled. It is what such a portfolio would plausibly contain:

- Government procurement at scale, converting a speculative revenue line into a sovereign one.
- Loan guarantees and public-private infrastructure financing.
- Tax incentives and accelerated depreciation for capacity investment.
- Power, transmission, and interconnection priority.
- Export controls and tariffs shielding domestic capacity.
- Restrictions on the deployment of foreign models in regulated or public-sector contexts.
- Strategic government equity or warrant positions.
- Regulatory protection, including compliance regimes whose fixed costs favor incumbents.

Each instrument is individually defensible. The composite is the conversion of private platforms into protected quasi-public infrastructure—the same administrative-convergence structure described in [Section 5.2](#), arriving through rescue rather than through compliance.

The relevant danger. Not that these firms are saved. That an open stack must then compete against subsidized incumbents whose compute, power priority, and regulatory position are underwritten by the sovereign, on the terms described in [Section 30.9](#). Capability audits will not distinguish the two; exit rights will.

I.4 Quarantined Claims

The source material for this appendix contains claims that are interesting and insufficiently supported. Listing them is not throat-clearing: an argument that adopts the useful half of a source and quietly discards the rest has not disclosed its own selection rule.

I.5 Indicators: What Would Confirm or Weaken This

Consistent with [Section 24.7](#), a scenario that cannot be checked is not worth writing. These are the observable series. They are grouped to match the loop, and none requires privileged access.

Treasury plumbing. Hedge-fund Treasury positions and their domicile; repo haircuts and funding spreads; futures basis and swap-spread positioning; dealer inventories; Treasury market depth during equity selloffs; whether long yields reverse upward after an initial risk-off

Claim	Why it is not used
A rescue of the AI sector has already been arranged	Asserted without evidence; unfalsifiable as stated
Open models were released specifically to impair a named foreign laboratory	Intent is not observable from release behavior; the competitive effect is what matters and is treated separately
AI accounts for 80–90% of recent U.S. GDP growth	No verifiable derivation; the loop in Appendix I.2 does not require it
Official gold revaluation to \$20,000–\$30,000 to extinguish debt	A policy speculation, not a stated or observed program
Imminent closure of the U.S. capital account	Contradicted by the stated posture examined in Section 30.3
Specific Treasury-yield “pain thresholds”	Precision unsupported by any published reaction function
A gold-based system would produce a 1–2% global risk-free rate Fed cuts, a bull steepener, mortgage-convexity buying, weaker dollar, stronger gold, and an equity rally will occur in that sequence	Insufficiently specified to evaluate A tactical path from commentary, not a state transition the thesis can assume; Section 27.2 uses quiet inflation compensation only as a tripwire for whether market repair remains available
Any named energy cutoff, maritime escalation, or infrastructure attack is a base-case forecast	The physical mechanisms are useful stress scenarios; event timing, intent, and joint probability are not established
Aggregate North American energy abundance guarantees local deliverability	Conversion, location, product, transport, and timing can still bind; DVC is designed precisely because aggregate stock is insufficient
An operator targeted sponsor interests intentionally because the effects harmed the sponsor	Harm and divergent incentives do not establish intent; only the principal–agent mechanism is used

Table I.1: Claims from the source material excluded from the thesis, with reasons.

decline; buybacks, issuance maturity shifts, and changes in regulatory treatment of sovereign paper.

AI financing. Hyperscaler free cash flow after capital expenditure; data-center project finance terms; vendor financing and lease obligations; private valuation growth *rates* rather than levels; credit-default swap spreads; margin loans collateralized by sector holdings; delayed rounds and down rounds.

Open-model economics. Price–performance rather than benchmark performance alone; open-weight adoption among enterprises; local inference and hardware compatibility; developer migration; whether proprietary providers are forced into persistent price cuts; restrictions on foreign model deployment.

Fiscal feedback. Non-withheld individual tax receipts; capital-gains and stock-compensation receipts; corporate receipts from technology firms; consumption sensitivity to equity declines; deficit behavior during even a modest correction; whether Treasury issuance rises while private sector borrowers seek capital simultaneously.

Industrial policy. Grid generation and transmission additions; transformer and switchgear capacity; machine-tool orders; skilled-trade wages and vacancies; semiconductor, shipbuilding, and critical-mineral subsidies; trade and investment restrictions.

Physical conversion and common cause. Refinery and product throughput rather than crude stock alone; firm-power availability; transformer and switchgear outages; cooling-water restrictions; hardware and firmware dependencies; cable and network-route failures; scenario-specific DVC, active minimum cuts, and substitution latency.

Settlement architecture. Official gold accumulation; use of non-dollar currencies in commodity settlement; dollar-stablecoin adoption abroad; tokenized Treasury growth; privacy and law-enforcement provisions in new payment systems; restrictions on self-custody or cross-border movement.

Native monetary buyers. Just-in-time fee acquisition; operator sell-through and inventory; burns and net issuance; self-custodied reserve accumulation; wrapper demand; leverage; holding-period distributions; loss-bearing capacity; and countercyclical accumulation during draw-downs and regime-pressure episodes.

Falsification. The collateral-loop framing is weakened if a significant equity drawdown passes through to receipts and deficits far more weakly than [Section 3.2](#) implies, or if a volatility shock produces the textbook flight-to-safety in long yields with no subsequent reversal. Both are observable within a single cycle. Record the prediction before the event, per [Section 24.7](#).

Appendix J

What Changed Across Releases

This appendix records what each version added, in the words used when it was added. It is release history, not argument: nothing here is required to read the thesis, and no claim depends on it. It is kept because a document that revises itself in public should say what it revised.

v1.2 additions: Layer 0 now includes **Facility Capacity Receipts** (Section 15.7) and **Physical VerifyPrice** (Section 20.8); Layer 6 telemetry now includes the **Wrapper Dominance Board** and **Agency Preservation Board** (Section 24.9). A new political-economy thread runs through Part I: the **Treasury Market as Control Panel** (Section 3.1.1), the **Participation Line** (Section 3.6), **Administrative Repression** (Section 5.2), and **AI Homestead vs. Enclosure** (Section 5.6), culminating in a ninth SoV requirement, **agency preservation** (Item 9), and the closing frame **Equip, Don't Manage** (Section 33.4).

v1.3 additions: A thread that runs *around* the stack rather than through it. The **Market Realization Plane** (Section 11.14) models how external financial wrappers—ETFs, treasury companies, custodians, derivatives, dealers, and systematic allocation rules—represent claims on the native asset, and **Corollary 11.1** establishes that price appreciation and monetary adoption imply each other in neither direction. This adds a fourth verification family, **VerifyFlow** (Section 24.8), a corresponding public board, a **compositional adversary** class that requires no hostile actor (Section 5.2), a third boundary tension between institutional holdability and native monetary function (Section 5.6), a split of allocator adoption into native (Phase II-A) and wrapper-led (Phase II-B) paths, **Market Realization Warnings** distinct from monetary red lines (Section 28.7), and the formal flow model in Appendix H. Readers who want only the conceptual payload can read Section 11.9.1 and Section 33.3 and skip the appendix entirely.

v1.4 additions: A thread that runs *beneath* the stack. Where v1.3 asked who captures the value, v1.4 asks what the substrate depends on. Layer 0 already instrumented physical fragility; it did not argue it. Three additions close that gap. **Energy & Physical Interdiction** (Section 5.5) adds an adversary class that raises the cost of verification without prohibiting anything—curtailment, tariff discrimination, interconnection denial, load prioritization—and separates the verification-side exposure (reference-hardware obtainability) from the proving-side exposure (bulk power), which behave differently and are routinely conflated. **Sovereign Optionality** (Section 15.8) aggregates existing Facility Capacity Receipt fields into a single published exposure, feeding risk haircuts so that issuance discipline extends from “we cannot see

this facility” to “we can see it clearly and it is fragile.” **Red Line 13** makes the dependency falsifiable: Physical VerifyPrice SLOs are exogenous to token price, as claimed, but not to the conditions under which power and hardware are obtainable—which means the hinge is externally triggerable, and should be monitored upstream rather than discovered at the hinge. [Section 15.10](#) resolves the resulting tension between demanding cheap verification and expensive redundancy via **disruption-adjusted VerifyPrice**, with resilience explicitly subordinate to the affordability hinge. Finally, [Chapter 30](#) adds the case the thesis previously argued only in the abstract: a state-scale **closed sovereign stack** that validates the Layer 0 premise while inverting who ends up sovereign, plus three objections it provokes ([Section 31.2.6](#), [Section 31.3.4](#), [Section 31.3.6](#)).

v1.5 additions: No new claims—this version makes the existing argument findable and harder to misread. An audit found that the thesis had accumulated excellent parts that the reader was left to assemble alone, so v1.5 adds a spine rather than a thread. The seven **premises** that had lived only on the project site are now stated in the document ([Section 1.2](#)), and [Section 7.1](#) walks the full conditional chain in one place: each of the seven links, where it is defended, and the condition that would sever it—closing with [Table 7.1](#), which rates our own argument link by link and marks two links *medium* rather than claiming uniform confidence. [Chapter 1](#) opens with reading paths so an allocator, a builder, or a skeptic can skip the cultural preamble and go straight to the load-bearing material. The Abstract and Executive Memo now say plainly what was previously implicit: that the hinge is externally triggerable, and that the competitor to an open stack is a competent closed one rather than fiat. [Chapter 31](#) is regrouped from a flat list into monetary, infrastructure, political, and market-structure blocks, so the strongest objection is no longer the twelfth item in an arbitrary sequence. [Table 22.3](#) maps each reference application to the layers it exercises, the boards it reports to, and the red lines it would trip—which incidentally shows that shipping applications cannot catch Red Lines 4 and 8. [Chapter 14](#) gains a physical-sovereignty question kept deliberately separate from the legibility question, because a facility can be perfectly transparent about being fragile. [Section 11.4](#) now walks all nine SoV requirements rather than the first seven, since the two it omitted—non-bypassability and agency preservation—are the two a merely useful capacity token would fail.

v1.6 additions: A dependency the thesis had described but not named. Where v1.4 asked what the substrate depends on, v1.6 asks what the *incumbent* order depends on—and finds a reflexive circuit rather than a hierarchy. The **collateralized sovereign stack** ([Section 3.2](#)) is the diagnosis: a state’s effective fiscal capacity now depends on the market value of the asset complexes it regulates, so asset prices, receipts, fiscal room, and Treasury functioning support each other recursively and fail together. [Section 3.1.2](#) identifies who actually clears the paper—leveraged, basis-sensitive intermediaries rather than patient foreign reserve managers—and separates **sovereign credit safety** from **collateral stability**, which are routinely conflated and

behave differently under stress. [Section 5.2](#) adds a fourth adversary class, *enclosure by rescue*: a state need not seize a firm it can simply protect, once that firm has become fiscally load-bearing. [Section 11.14.2](#) extends the Market Realization Plane with **numeraire-dependence**, since a measured price path is jointly determined by the asset and the unit it is measured in. [Section 30.3](#) reports that the doctrine of [Chapter 30](#) is now arriving in the West as well, so the thesis faces two state-key stacks rather than one, and [Section 30.4](#) separates open weights from open sovereignty—released weights confer neither independent power, hardware, data, communications, privacy, nor settlement. Material that could not be verified to the standard the body requires is quarantined in [Appendix I](#), which is labelled non-load-bearing and lists what was excluded and why. Nothing was renumbered: Layers 0–6 are unchanged, and no red line or value-capture condition depends on any macro claim added here.

v1.7 additions — the fee-level argument and a relocated monetary claim. This entry and the one that follows it record subtractions rather than additions, and both landed after the v1.6 artifact was frozen—a reader checking the v1.6 PDF will not find them there. The [Value Capture Lemma](#)’s five conditions all govern *where* fees go; none governs *how large* they can be. [Section 11.10](#) closes that gap and does not like all of the answer. A protocol fee is a wedge on turnover rather than a claim on residual profit, so a contestable market can be taxed and operator margin compression is irrelevant to fee revenue—but the sustainable fee is bounded by the marginal buyer’s willingness to pay for the protocol’s differential properties, net of the fact that verified compute costs *more* to produce, and that willingness is probably thin wherever a creditworthy, suable indemnitor substitutes for a proof. Bypass is accordingly a magnitude rather than a binary, and two bypass channels were not being counted: partial bypass at the intensive margin, and a fork of the protocol with a lower fee and no burn. The consequence is that fees, burns, and collateral establish a competitively priced cash-flow claim and a balance-sheet floor—both discounted-cash-flow quantities—and not a monetary premium. [Section 11.10.8](#) relocates the monetary claim to the state-contingency of that differential value, separating two mechanisms that had been run together: the countercyclicality of the fee stream, which is a negative-beta discount-rate effect a discounted-cash-flow valuation captures exactly and which is therefore conceded to confer no moneyiness; and a **holder-side service flow**, which is not a distribution at all and so has no stream to discount. The second is where the monetary claim now lives. It is then discounted honestly: the states where the hedge is most valuable are the states where the protocol’s ability to supply is most impaired. The mechanism is standard; its magnitude is unsized and is presented as unsized. [Section 11.10.9](#) shows collateral is unit-elastic in price, bounded by an accounting identity to large-cap-equity magnitudes, and subject to wrong-way risk no red line previously covered. Two chain-strength self-ratings were reduced accordingly ([Table 7.1](#)), and **Red Line 14** makes the fee-level question falsifiable, bringing the published count to fourteen.

Also in v1.7 — the Bitcoin objection, and a correction to how this document talked about

proof-of-work. The objections chapter answered gold and did not answer Bitcoin, which for a store-of-value thesis is the more consequential omission. [Section 31.1.1](#) now opens the monetary block and concedes more than is comfortable: Bitcoin already satisfies the requirements set out here, on a fifteen-year record, through a work function with *no external buyer* — and that absent buyer is the source of its objectivity, not an inefficiency. Proof-of-useful-work weakens it. The thesis’s actual claim is narrower than earlier drafts implied: not that a better money is available, but that Bitcoin does not by itself supply portable attestation or verified compute, that these are separate goods with separate demand, and that whether they carry monetary premium or price as services is open and measurable. Passages in Parts I, II, III, V, and VI that described SHA-256 work as “waste” or “heat” have been rewritten to state the trade honestly—**monetary objectivity exchanged for capacity relevance**—and a succession framing running from Part I to the conclusion has been made non-hierarchical, since the thesis claims a different function rather than a later stage. What losing to Bitcoin would look like is specified as **Service-Good Realization** ([Section 28.7](#)), and is deliberately filed as a market realization warning rather than a fifteenth red line: its price-comparative clause cannot be a red line without breaking [Corollary 11.1](#). The count stays at fourteen.

v1.8 additions — duration of the claim, duration of the project. A completeness amendment, not an architectural one. Layers 0–6 are not renumbered. [Section 3.3](#) distinguishes duration of the *claim* from duration of the *project*: a repression-resistant store of value must remain duration-neutral, while civilization still requires a duration warehouse capable of holding the construction interval of plants, grids, and data centers. The two must not be the same instrument. [Section 3.1.2](#) extends the residual-buyer account from leveraged intermediaries to automatic index flows that absorb less DV01 per dollar after a selloff. [Chapter 30](#) names China’s directed banks as a duration warehouse and the Western bill-heavy book as sovereign maturity transformation. VerifyFlow gains a buyer-quality checklist, not a scalar index ([Section 24.8.7](#)). A new objection, [Section 31.5.1](#), concedes that neutral money does not finance the reactor and refuses to let proofs, Work Credits, or the native asset become a coupon by another name. Tactical bond-market claims are quarantined in [Appendix I.4](#).

v1.9 additions — deliverable service, monetary risk absorption, and architectural agnosticism. The core claim is narrowed: a bearer base asset may earn monetary premium only where the full service path remains stress-deliverable, native demand cannot bypass it, a persistent self-custodied loss-bearing constituency absorbs residual risk, and project credit remains separate. The steel-man expands to ten premises and the conditional chain to nine links. [Section 7.2](#) adds the Topological Scarcity Lemma, the inverted-U Pressure–Capacity Corridor, and the Triad Coherence Test comparing one-token, neutral-reserve-plus-credit, and shared-settlement-plus-modular-collateral designs. [Section 15.9](#) restores conversion throughput through scenario max-flow/min-cut, substitution latency, and common-cause dependency analysis; issuance is bounded by stress-adjusted DVC rather than nameplate capacity. [Sec-](#)

tion 24.8.8 moves buyer quality into the monetary chain, while the public boards add common-cause and sponsor dependencies. Section 27.2 separates four policy states without turning them into tactical forecasts. Section 28.5 compares three collateral models, and Red Line 15 makes the native collateral–capacity spiral measurable. Section 31.5.2 specifies explicit proof-audited project credit with a constitutional prohibition on transferring par, redemption, or emergency-support promises to the base asset. Work Credits are reclassified as typed service claims; FCR, FER, and PIDL artifacts as evidence; project notes as duration-bearing credit; and LP/staking positions as derivatives or operating claims.

v2.0 — the assessment pass. No new claims: this version makes the v1.9 text agree with its own falsifiability standards. A systematic adversarial review resolved five internal contradictions. Gold-tier sampling coverage is counted in production lots at tier minimum rates, with the rule-of-three converse stated (Section 15.4). A maturity gate blocks Tier A (pristine-collateral) issuance while any PoUW verification component is Experimental or Pilot—currently binding on every PoUW workload (Section 20.10). The absolute $\varepsilon = 10^{-5}$ bound on FP32 matmul verification is replaced by a scaled relative bound with a pinned reference kernel, and “transcript determinism” is restated as pinned-kernel reproducibility (Section 20.3). The Red Line 8 emergency path is reconciled as subtractive-only, with routine use of subtractive powers tripping the red line and signer accountability specified (Section 23.5.3). Naive demand-to-price passages in Part II are explicitly marked as routing rules rather than findings, with the refutation cross-referenced where it lands (Section 11.9). Four analytical gaps are given measurement machinery: a fifth measurement family, convenience-yield telemetry, prices the holder-side service flow through base-asset lending rates, forward and perpetual basis, and wrapper basis (Section 24.8.6); the joint failure mode—closed-stack services plus wrapper price realization leaving the asset a reference price with no native loop—is named with a triple-divergence diagnostic (Section 30.9); stablecoins are engaged as the incumbent bypass channel (Section 11.12.1); and the utility-token record is cited against the thesis, with the 30% fee-coverage gate labeled an order of magnitude above the precedents (Section 31.1.4). Legal posture is narrowed: a securities-law analysis confronts the Howey structure of the value-capture design itself (Section 25.2), and the AML table is reframed as requiring legislative change rather than as a boundary regulators must respect, with the Tornado Cash precedent added to defense-in-depth (Section 25.1). Verification expectations are aligned with engineering reality: INFER_LM_70B is re-tiered to aspirational with no Sev-1 threshold until ZKML maturity lifts the gate, and Phase I expects INFER_LM_7B (Section 27.3). The final chapter is retitled “Why These Could Become Money,” and the base-rate problem is stated: no asset in monetary history is known to have crossed from utility cash-flow claim to monetary premium (Chapter 32). Prior-art acknowledgments are added for Layers 0–2 (TUF/Uptane, Certificate Transparency, OONI, Tor/I2P, BIP-324, BOLT12, energy resource-adequacy literature) with matching Sources entries. The simultaneity problem in elasticity estimation—wrapper

and index flows moving NAV—is addressed in the VerifyFlow Measurement Contract ([Appendix H](#)). No red line was relaxed, no threshold was moved, and no falsification condition was withdrawn.

v2.1 — the McLuhan pass. No new claims: this version adds a framing layer, not evidence. The thesis’s arguments are placed in the lineage of media-transition theory (McLuhan 1962, 1964; Eisenstein 1979; Ong 1982; Anderson 1983; Sources), with every use explicitly bounded as mechanism rather than measurement. Eight placements, none load-bearing. The Executive Memo gains a closing paragraph explaining why the thesis instruments rather than forecasts: every prior restructuring of the information substrate arrived as an externality of adoption, and the consequences were nobody’s plan. [Section 2.0.1](#) adds a prior-transitions passage to [Chapter 2](#): manuscript trust anchored in persons, print in artifacts, broadcast in institutions, and the networked era arriving with no native anchor at all—the slot this thesis proposes to fill with proof. [Chapter 8](#) gains an interiority paragraph: the private self as a print externality (Ong), which reframes surveillance as the withdrawal of a centuries-old subsidy rather than the violation of a natural constant. [Chapter 3](#) opens with the topology observation: truth and value move on the same topology, and both are moving from broadcast ($1 \rightarrow N$) to network ($N \leftrightarrow N$); the state as typesetter of uniform value belongs here. The canonical-workload discussion ([Section 20.3](#)) names its precedent: typographic standardization preceded and enabled interchangeable parts (Eisenstein), with the necessity claim explicitly not adopted. The Wrapper Dominance Ratio ([Section 11.13](#)) gains a reading rule: first-generation use of a new medium wears the old medium’s form, so a rising WDR measures a phase, and the failure gate exists because the frame does not promise completion. [Chapter 32](#) tacks two sentences onto the base-rate paragraph: a printed scroll is still a scroll; the precedents shipped content without structure, which is why the base rate is empty and why the empty base rate is not dispositive—and the reading earns nothing, because the chain still has to prove itself in telemetry. The VerifyPrice SLO sidebar ([Section 20.7](#)) gains the habituation argument: if a medium restructures only those who habituate to it, verification must be cheap enough to become a habit before it can become anything else, which is why the exogeneity requirement is the precondition of the medium having any effect at all. Six Sources entries are added with scope notes stating what each is and is not cited for; the “we shape our tools” line is attributed to Culkin (1967), not McLuhan. No red line was relaxed, no threshold was moved, no falsification condition was withdrawn, and no claim of the thesis now depends on the frame.

v2.2 — the instrument-hardening pass. No new claims: this version makes the two flagship measurement families agree with the standards the rest of the thesis already holds itself to. An audit of VERIFYPRICE and VerifyFlow found the core derivations sound—every equation of [Appendix H](#) re-derives and the recycling boundary verifies numerically—but the aggregation and specification layer behind them had ten defects, now fixed. **Aggregation:** the asset-level vector’s $Q^{\text{exp}} = Q^{RC} + Q^{RD}$ double-counted the ε -driven creation flow, which sits inside κ

in Q^{RC} and again inside raw net creation in Q^{RD} ; Q^{RD} now consumes only the intercept component, per the decomposition rule [Appendix H](#) already stated but the vector violated, with the regression residual published as its own series and consumed by no aggregate ([Section 24.8.2](#)). The Wrapper–Native Growth Gap, which feeds the triple-divergence composite and Red Line 9, gains the aggregation specification and deflation adjustment its adjacent Red Line 14 quantity already carried; an unadjusted WNG can rise on a valuation artifact alone ([Section 24.8.2](#)). The Mechanical Pressure Ratio now publishes gross and net forms, mirroring the Wrapper Dominance Ratio’s stock/flow split ([Section 24.8.2](#)). **Domain:** the Wrapper Recycling Ratio is undefined at $L = 1$ —the unlevered wrapper class most likely to dominate this asset—so the κ –WRR identity is stated, the $L=1$ publication rule fixed (publish ε ; never a WRR column), and a wrapper-class applicability map added so classes without shares and NAV (perpetual futures above all) get named substitute constructions instead of fabricated readings ([Appendix H.5](#), [Appendix H.6](#)). **Anti-gaming:** the $r(W)$ production denominator gains the measurement contract its numerator always had—reference cost models, pinned like the reference kernel, independently estimated—because a prover inflating production cost drives the ratio down and the network looks more verification-asymmetric than it is ([Appendix A](#)). **Estimation:** the OLS-as-upper-bound rule for flow elasticity was conditional, not absolute: for a stale-priced underlying, errors-in-variables attenuation biases the estimate downward, and an attenuated OLS value bounds nothing; a staleness diagnostic is now published alongside, and where staleness is material the reading is published as unidentified ([Appendix H.1](#)). **Specification:** VerifyPrice carried three tuple definitions across four sites; [Appendix A](#) now states the canonical five-field form and demotes the others to views of it. VerifyPower took the reciprocal of a vector; it is now a scalar on a pinned reference workload. VerifyPrice_{adj} is annotated as a convenience summary no red line reads on, with the per-state vector load-bearing, because an expectation averages away the tail that resilience lives in ([Section 15.10.3](#)). The AffordabilityRatio’s perverse direction—improving when fees rise—is named, with denominator floors and a USD-referenced settlement variant ([Section 20.7](#)). Every fix tightens an instrument; none relaxes one. No red line was relaxed, no threshold was moved, and no falsification condition was withdrawn.

v2.3 — the coherence pass. No new claims: a systematic audit of self-coherence, local coherence, and global coherence produced 125 severity-ranked findings, all remediated. **Two mathematical errors:** the USD price model multiplied Joules directly by a USD/kWh price — a dimensional error overstating the energy term by 3.6×10^6 — now derived per run with the explicit conversion and cost percentiles computed as percentiles of per-run cost, never recomputed from percentile inputs ([Section 20.6](#)); and the WNG revaluation artifact was stated backwards — the phantom arises on a price *rise* (wrapper exposure is a stock revalued with every tick; native use accumulates at transaction prices), now stripped by unit-denominated aggregation ([Section 24.8.2](#)). **Conventions pinned once:** Q^{RD} owns the intercept α only, with

the regression residual published as its own series and consumed by no aggregate; the take-rate τ uses gross fees while FeeCoverage uses retained fees, dissolving every fee-plus-burn ratio and retiring the label; the WDR is a pair and Red Line 9 requires both components to rise; two distinct deflators — WNG unit-denomination versus Red Line 14-B’s workload unit-price index — are named and kept separate. **Red lines:** RL14-B gains persistence and a named deflator; RL11/RL13 thresholds instantiated where “e.g.” placeholders stood; RL15’s clauses labeled conjunctive; severity labels (“kills” vs. “weakens”) defined; response-protocol clocks given pinned start triggers and a no-reset rule. **Instruments:** VerifyPower carries the token price explicitly ($q_{\text{fee}} \cdot P_X / p_{50,c}$); the AffordabilityRatio’s undefined “floor” replaced by three concrete guards; the Part III decision rule gates on the worst disruption state’s per-state bound with the π -weighted average demoted to tie-breaker; the canonical tuple reconciled across all five sites. Roughly forty definitional repairs across Parts I–VI, the appendices, and the executive memo, including the Value Capture Lemma’s failure clause (no *monetary premium*, not no value at all), the Work-Credit supply identity, the θ_t double-count, the φ_t false equivalence, and the base-rate defense restated in its honest two-move form. No red line was relaxed, no threshold was moved, and no falsification condition was withdrawn; several conditions became checkable that previously were not.

v2.4 — the household-transmission pass. New claims, all bounded: this version adds the household layer to a thesis whose adversaries had been exclusively sovereign and financial. The source is Green’s “The Vibecession Was Real” (Sources), adopted for mechanism and quarantined for quantities, simulations, and psychology. **The social contract as a failing soft guarantee:** the behavioral promise that discipline converts into advancement — honored contract by contract while failing in aggregate, because the incumbent’s balance-sheet wealth is the entrant’s acquisition price — is added to Chapter 2 as a trust failure that survives audit. **The exhaustion gradient:** the participation line (Section 3.6) gains its approach path — formally inside the system while progressively less able to accumulate — with the observation that dollar-weighted aggregates lag person-weighted depletion until a delinquency threshold is crossed. **Capability versus formal access:** the ninth SoV requirement (Item 9) gains its second layer: a valid balance is not command over the necessities of an independent life, and a ledger can supply the first while failing the second. **Representation discipline:** every board aggregate now discloses its weighting rule, because dollar-weighted and person-weighted summaries of the same observations can move in opposite directions ($\bar{x}_{\text{dollar}} - \bar{x}_{\text{person}} = \text{Cov}(C_i, x_i) / \bar{C}$); neither is the false one, and where they diverge persistently both are published (Section 24.8.8). **Household transmission in the state model:** Section 27.2 gains cross-cutting observables — saving rate, delinquency by cohort, first-time-buyer share, underemployment, mean-versus-median net worth, sentiment against spending — with the argument that household depletion changes how states transition, not just how they feel. **Distributed resilience:** Section 30.9 gains the formal model ($O_i < 0 \Rightarrow \Delta R_i < 0$; $R_{\text{state}} \neq R_{\text{household}}$) and the requirement that the

open stack evidence distributed resilience rather than sovereign continuity alone — a stability argument, not a redistribution argument. **The household amplifier:** the joint failure mode paragraph gains the nonlinear-break mechanism — resilient aggregates as the sum of two individually fragile mechanisms, upper-cohort discretionary spending and lower-cohort necessity spending, each of which fails differently under the same asset repricing. **A new objection** (Section 31.3.5): a perfect ledger hardens an unequal starting state; the response concedes the ledger point, locates the real claim in acquisition-channel breadth rather than record neutrality, and specifies the falsifier — new-participant acquisition from labor income staying negligible while incumbent concentration rises, visible in buyer-map cohorts. Six glossary entries and the synthesis’s closing edge: proof tells us whether a claim is true, representation whom it describes, capability what participants can do — a stack that verifies everything and represents almost no one is a more efficient ledger for the exhaustion dynamics it was built to escape.

v2.5 — the prose pass. No new claims: this version makes the argument read like itself. An audit against the thesis’s own best prose (Part I) found the gold-standard voice allocated almost entirely to the demolitions (fee incidence, the fork, the closed-stack comparison) while the constructive chapters (Work Credits, the design space, the evaluation framework) and the instrumentation core (telemetry, red lines 1–8) were left as bullet lists and spec boxes. The capability was present throughout — the audit found Part I-grade passages inside nearly every weak section — but the document’s spec-first architecture kept burying them. **The red lines get their paragraphs:** Red Lines 1–8, previously bare definitionboxes with one-line shrugs, now carry the dramatization Red Lines 9–15 already had — the touchstone for RL1, refund safety as a verdict for RL2, monoculture as a single point of political failure for RL3, telemetry capture as rigging the scoreboard for RL4, the guards-in-scrip condition promoted to trip wire for RL5, the paper-gold precedent for RL6, the shrinking anonymity set for RL7, and precedent-not-heist for RL8. **The conclusion is re-ordered and its overclaim fixed:** the “Proofs will anchor truth” formulation — explicitly disavowed in Section 31.2.3 (“proofs do not create truth; they bound disputes”) — becomes “Proofs will bound what can be faked”; the seven-instrument incantation at Section 33.4’s peak is replaced with three human sentences each instrument makes measurable; and “Two Loops, Not One” now precedes “Equip, Don’t Manage,” so the thesis ends on the owner/tenant question rather than a dashboard epigraph. **The reference applications become narratives:** the four Create/Compute/Prove/Settle/Verify recitations (Chapter 22) now open with a scene — a newsroom and a fabricated clip, a treasurer under capital controls — before the flow. **Analogies deployed throughout, all historically on-thesis:** the touchstone (verification asymmetry), Lloyd’s Register and the assay punch (lot sampling), water pressure and the trunk main (DVC and min-cut), paper gold (wrapper dominance), the fire extinguisher versus the contract for its delivery (convenience yield), SushiSwap’s vampire attack (the fork’s moat), a thousand cars with one GPS (policy concentration), guards paid in fresh scrip (issuance-funded security), the toll versus the detour (the capturable wedge). **Struc-**

tural repairs: the Part III closing refrain, printed twice within fifty lines (the L3 close and the Part III summary), is deduplicated and the summary given its own close; the 230-word Mechanical Pressure Ratio bullet is broken into prose; the “Cross-reference:” ledger-opener in fee routing is replaced; the hedging loop (value-capture conditions restated four times, “typed service claims” recited) is compressed to statement-once-plus-pointer. No mathematical content, definition, threshold, or cross-reference was altered; every ??arget and urvives.