

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 1.5

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 1.5. Licensed CC BY 4.0.

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; synthetic media and platform moderation undermine the assumption that “seeing is believing”; AI concentrates cognition and compute behind hyperscaler 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 makes a concrete claim: **three cryptographic capacities can function as monetary primitives** for a dense digital civilization: **Privacy** (censorship-resistant settlement that preserves agency), **Proofs** (portable attestations of computation and provenance), and **Compute** (useful work such as matmul, inference, and ZK proving, all wrapped in succinct guarantees). They become monetary only if they preserve agency, remain cheap to verify, survive repression, and convert recurring demand into scarce, non-bypassable asset value through fees, burns, collateral, disciplined issuance, and public telemetry. 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 **VerifyReach**, **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 mea-

sure 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. A triad asset becomes money only when Privacy, Proofs, and Compute generate recurring native fees, burns, collateral, settlement, and verified-work demand under healthy SLOs; 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. 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—conditions partly set by states through curtailment, rationing, tariff discrimination, interconnection denial, and load prioritization. The hinge is therefore externally triggerable. We separate the verification-side exposure (reference-hardware obtainability) from the proving-side exposure (bulk power), aggregate the substrate risk as **sovereign optionality** over Facility Capacity Receipt fields, and bind it to issuance through risk haircuts and a dedicated red line. 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. Capability was never the claim; the properties are. In total the thesis publishes **thirteen 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. 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. The stack becomes monetary only if the asset captures indispensable demand and cannot be bypassed by fiat, stablecoins, cloud contracts, custodial wrappers, or administrative convenience.

The thesis is conditional. Gold and hard assets are the bridge. Privacy, Proofs, and Compute are the candidate destination. Telemetry decides when the bridge can bear less weight.

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.

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

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 **thirteen red lines** rather than promising success.

Contents

Abstract	ii
Executive Memo: Why AfterFiat Exists	iv
Reader’s Map	xvi
Thesis in Plain Language	xxii
Key Definitions	xxiv
I Context & Claims	1
1 Introduction	3
1.1 Central Claims	6
1.2 Premises	8
1.3 Contributions	8
1.4 End-to-End Vignette: The Loop in Action	10
2 The Failure of Soft Guarantees	14
3 The World Forces New Monetary Primitives	16
3.1 Macro Playbook: Debt Repression Flight to Neutrality	16
3.1.1 The Treasury Market as Control Panel	17
3.2 The Repression Playbook, Then and Now	18
3.3 Social: The Web’s Trust Default Has Flipped	19
3.4 Political: The Participation Line Has Broken	21
4 First Principles: What a SoV Must Survive	22
5 Threat Model	26
5.1 Scope and Assets	26
5.2 Adversary Classes	27
5.3 Threats by Loop Stage	30
5.4 Mapping the Authoritarian Playbook	31
5.5 Energy & Physical Interdiction	32

5.6	Explicit Tensions and Boundary Conditions	34
6	Layers of the Cypherpunk Stack	38
6.1	Modules in Outline: Applications and Primitives	41
II	The Triad as Monetary Base	43
7	The Triad as Monetary Base	45
7.1	The Conditional Chain, End to End	46
7.2	Monetary Objects and Value Capture	48
7.3	Three Reference Designs	51
8	Privacy as Private Money	53
9	Proofs as Attestation Money	55
10	Compute as AI Money	57
10.1	Compute Deflation and What Remains Scarce	57
11	Work Credits: Energy-Anchored Claims	60
11.1	Definition: Work Receipts vs. Work Credits	60
11.2	Energy Anchoring	62
11.2.1	Layer 0 Maturity and Economic Consequences	62
11.3	Robots, AI, and the Demand for Work Money	63
11.4	Why This Is Still “Money,” Not Just Coupons	64
11.5	Monetary Role	65
11.6	Issuance: Tying Credits to Real Capacity	65
11.6.1	The Supply Balance Equation	66
11.7	Energy-Priced, Not Energy-Pegged	67
11.8	Why Work Credits Are SoV, Not Just a Utility Token	68
11.9	Economic Linkage: How Triad Demand Becomes Asset Value	70
11.9.1	Value Capture vs. Price Capture	73
11.10	From Utility Demand to Monetary Premium	73
11.11	Why Users Cannot Simply Bypass the Asset	74
11.12	Wrapper Dominance Risk	75
11.13	The Market Realization Plane	76
11.13.1	A Three-Stage Model	77

12 Monetary Design Space	79
12.1 ZK Money	79
12.2 AI Money	79
12.3 Hybrid Instruments	79
13 ZK Money and the ZK + AI Economy	81
13.1 Why Open Hardware Is a Precondition for the ZK-Economy	81
13.2 How This Opens a ZK + AI Economy	82
14 SoV Evaluation Framework	84
14.0.1 Comparison: Traditional Hard Assets vs. the Triad	86
III Infrastructure Layers 0–3	88
15 Layer 0: Verifiable Machines & Energy	90
15.1 Why Layer 0 Is a Monetary Question	90
15.2 Design Goals and Non-Goals	91
15.3 Hardware as Base Reality for Work Credits	92
15.4 The Layer 0 Feasibility Ladder	92
15.4.1 Policy Hooks: Economic Treatment by Grade	93
15.4.2 Quantitative Thresholds (Reference Design)	94
15.4.3 Pragmatic Starting Point (L0-A)	95
15.4.4 The Path Forward	95
15.4.5 Why This Prevents “Layer 0 Is Impossible, Therefore Thesis Fails”	95
15.5 Concrete Components of Layer 0	96
15.5.1 Open Designs Where Possible	96
15.5.2 Attested Randomness and Entropy	96
15.5.3 Attestation Without Priesthood	96
15.5.4 Lot Sampling and Destructive Audits	96
15.5.5 How This Opens New Economies	98
15.6 Verifiable Power: Energy as First-Class Input	99
15.7 Facility Capacity Receipts: Beyond Energy	99
15.8 Sovereign Optionality: Aggregating FCR Into One Exposure	100
15.8.1 Adaptation of Scale	101
15.8.2 Mapping to Existing FCR Fields	101
15.8.3 Facility and Network Form	101
15.8.4 Governance of the Weights	102
15.8.5 What $\mathcal{O}_s$ Is For	102
15.9 Resilience Versus Affordability: The Layer 0 Cost Tension	103

15.9.1 Two Different Exposures	103
15.9.2 Disruption-Adjusted VerifyPrice	104
15.9.3 The Decision Rule	105
15.10 Operational Patterns: Profiles, Upgrades, and Failure Modes	105
15.10.1 Hardware Profiles and Workload Binding	105
15.10.2 Upgrades and Deprecations	106
15.10.3 PQ and Cryptographic Agility	107
15.11 Stress Tests for Layer 0	107
15.12 What Layer 0 Exports to Higher Layers	107
16 Layer 1: Reachability	108
16.1 Threat Model: What We Must Survive	108
16.2 Design Rules: Protocol Posture Under Pressure	108
16.3 Mechanisms: What Needs to Ship	109
16.4 How Layer 1 Anchors the Triad	109
16.5 VerifyReach: Communications Telemetry	110
16.5.1 VerifyReach Measurement Specification	111
17 Layer 2: Distribution & Execution	113
17.1 Threat Model & Objectives	113
17.2 Design Rules	113
17.3 Reference Distribution Architecture	114
17.4 Key Management & Release Process	115
17.5 Update Economics & Neutrality	115
17.6 Fallback Playbooks	115
17.7 Lawful Privacy in Distribution	116
17.8 Update Telemetry	116
18 Layer 3: Identity & Claims	117
18.1 Why Identity Must Decouple from Doxxing	117
18.2 Design Goals	117
18.2.1 Issuer Pluralism Without Issuer Anarchy	118
18.3 The Identity Kernel	119
18.4 Human Identity: Personhood and Compliance Without Dossiers	119
18.4.1 Sybil Resistance Menu	120
18.5 Machine Identity: Capability, Not Brand	121
18.6 Reputation: Priced Behavior, Not Personal Data	121
18.7 Patterns That Keep Identity Private and Usable	121
18.8 Failure Modes & Guardrails	122

18.9 How Identity Fits the Modular Stack	122
18.10 Identity/Reputation SLOs	122
18.11 Implementation Sketch	123
18.12 Part III Summary: Minimum Viable Stack	123
 IV Truth, Work, and Settlement: Layers 4–5	 125
 19 From Infrastructure to Economics	 127
 20 Layer 4: Truth & Work	 128
20.1 What Layer 4 Is (and Isn't)	128
20.2 Verification Asymmetry Revisited	128
20.3 Canonical Workloads and Proof Types	129
20.3.1 Canonical Workload Definition Template	129
20.4 PoUW Design Patterns	131
20.5 Proof Factories and PaL	132
20.6 VerifyPrice Observatory	133
20.6.1 VerifyPrice Measurement Specification	133
20.7 VerifyPrice in Practice	135
20.7.1 Physical VerifyPrice SLOs (Constitutional)	135
20.7.2 Protocol Affordability SLOs (Operational)	136
20.7.3 Token-Quoted VerifyPrice (Market Signal, Not Target)	136
20.8 Physical VerifyPrice: Auditing the Infrastructure Behind a Workload	136
20.9 Layer-4 Stress Tests	137
20.10 Minimum Viable Layer-4 Economy	137
 21 Layer 5: Value & Settlement	 139
21.1 Settlement as a First-Class Workload	139
21.2 Design Constraints for Privacy Rails	140
21.3 The Privacy Rails Kit (PRK)	140
21.4 Settlement Safety: Refund and Bridge Invariants	141
21.5 VerifySettle: Settlement Telemetry	141
21.6 Layer-5 Stress Tests	142
21.7 Minimum Viable Layer-5 Corridor	142
 22 The Modular Stack	 144
22.1 Primitive Catalog: The Twelve Pieces	144
22.2 The Four Reference Applications at a Glance	146
22.3 Reference Application: Private Treasury & Payroll	146

22.4	Reference Application: Media Provenance	148
22.5	Reference Application: Verified Inference	148
22.6	Reference Application: Proof/Compute Procurement	148
22.7	Part IV Summary	149
V	Governance, Telemetry & Neutrality	150
23	Layer 6: Governance & Telemetry	152
23.1	“No Dashboards, No Trust”: Public SLOs as Constitution	152
23.2	Control Surfaces: Parameters, Upgrades, Emergency Powers	153
23.3	Bell-Labs-Style R&D vs. On-Chain Governance vs. Off-Chain Norms	153
23.4	Governance as SLOs: The Five Invariants	154
23.5	Constitutional Enforcement	157
23.5.1	Machine-Enforced Constraints (Hard)	157
23.5.2	Override Path (Slow, Expensive)	157
23.5.3	Emergency Path (Narrow)	157
23.6	Monetary Constitution	158
23.6.1	Issuance Envelope	158
23.6.2	Fee Routing	158
23.6.3	Retirement Rule	158
23.6.4	Risk Haircuts	158
23.6.5	Coverage Target	159
23.6.6	Value Capture Loop	161
24	Extended Telemetry	162
24.1	Reference Verifier Classes	162
24.2	Reference Verifier Implementations	162
24.3	Cost Vector Definition	163
24.4	Sampling and Anti-Cherry-Pick	163
24.5	Data Availability and Anti-Memory-Hole	163
24.6	Reproducibility Contract	164
24.7	How to Measure: Receipts, Not Vibes	164
24.8	VerifyFlow: Instrumenting the Market Realization Plane	165
24.8.1	Wrapper-Level Vector	166
24.8.2	Asset-Level Vector	166
24.8.3	Holder Cohorts, Vintage, and Policy Concentration	167
24.8.4	Capital Survival: AUM Is Not Investor Return	167
24.8.5	Dealer Balance Sheets and Hedge Paths	168

24.9 The Public Boards	169
24.10 Detecting Drift	170
24.11 Making Neutrality and Repression-Resilience Falsifiable	170
24.12 Policy-Attack Stress Harness	171
24.12.1 Market-Realization Stress Tests	171
25 Legal, Policy, and Jurisdictional Posture	173
25.1 Lawful Privacy as Protocol Design	173
25.1.1 What Lawful Privacy Is Not	173
25.1.2 Social Coercion Metrics: When “Optional” Becomes Mandatory	174
25.2 Defense-in-Depth Against Bans and Sanctions	174
25.3 Labs, Foundations, and Neutral Router Commitments	175
26 Operator & Investor Checklist	176
26.1 Triad Supply: Privacy, Proofs, Compute	176
26.2 Value Capture & Anti-Bypass	176
26.3 Institutional Holdability	177
26.4 Stack Health: Layers 0–6	177
26.5 Telemetry Honesty	178
26.6 Red Flags and Failure Patterns	178
VI Dynamics, Risk & Implementation	180
27 Adoption Curve & Ecosystem Dynamics	182
27.1 The Transition Window: Hard Assets as Bridge	182
27.1.1 Portfolio Migration: From Bridge Assets to Triad Assets	183
27.2 Four Phases of Adoption	183
27.2.1 Phase Gates and Failure Gates	184
27.2.2 Adoption Curve To Expect	185
27.2.3 Phase I: Cypherpunk-Led (Prove the Plumbing in Public)	185
27.2.4 Phase II: Allocator-Led (Proofs and Privacy Become Budget Lines)	186
27.2.5 Phase III: Composability-Led (The Stack Disappears into Infrastructure)	188
27.2.6 Phase IV: Policy-Led and Path-Dependent (Verification as Public Good)	189
27.2.7 The Negative Case: Useful Infrastructure Without Money	190
27.3 Failure Gates	191
28 Risk Analysis & Failure Modes	192
28.1 Technical Risks	192
28.2 Economic Risks	192

28.3	Political Risks	192
28.4	Spec Drift, Vaporware, and Scoreboard Capture	193
28.4.1	Scoreboard Capture in the Market Realization Plane	193
28.5	Red Lines: When the SoV Thesis Fails	194
28.6	Market Realization Warnings	197
29	Implementation Sketches for Builders	199
29.1	Minimum Viable Stack (MVS) Sequencing	199
29.2	Layer-0 Verifiable Machines	200
29.3	Privacy Rails: Making Settlement Safe and Boring	200
29.4	Proof Factories: Receipts as a Service	200
29.5	Compute Consensus Pilots	201
29.6	Developer Kit: Make “Import Proofs” the Default	202
30	The Closed Sovereign Stack	203
30.1	A State That Already Assumes Soft Guarantees Have Failed	203
30.2	Trust Minimization Applied to Matter	204
30.3	The Inversion: Who Is Sovereign	205
30.4	What This Validates	206
30.5	What This Warns	206
30.6	What We Do Not Claim	207
30.7	What This Changes in the Stack	208
31	Objections & Responses	209
31.1	Monetary Objections	209
31.1.1	“This is just another chain / coin.”	209
31.1.2	“Utility does not imply money.”	210
31.1.3	“The asset can be bypassed.”	210
31.1.4	“Gold still beats this because gold needs no network.”	210
31.2	Technical and Infrastructure Objections	211
31.2.1	“Users don’t care about privacy or proofs.”	211
31.2.2	“PoUW will centralize / can’t compete with hyperscalers.”	211
31.2.3	“Proofs do not prove truth.”	212
31.2.4	“They’ll just shut off the internet / app stores.”	212
31.2.5	“This burns too much energy.”	212
31.2.6	“You will lose the competition for electrons.”	213
31.3	State, Political, and Governance Objections	213
31.3.1	“Governments will never allow lawful privacy at scale.”	214
31.3.2	“Receipts will become surveillance tools.”	214

31.3.3	“This just gives the state better compliance tools.”	215
31.3.4	“The closed sovereign stack simply wins.”	215
31.3.5	“This is political preference dressed as analysis.”	216
31.3.6	“AI will be enclosed anyway.”	216
31.3.7	“Governance will just re-centralize.”	217
31.4	Market-Structure and Co-option Objections	217
31.4.1	“The asset will be co-opted, not bypassed.”	217
31.4.2	“This will just become another ETF.”	218
31.4.3	“The price went up, so the monetary thesis is working.”	218
31.4.4	“Wrappers increase access, so co-option is harmless.”	218
31.4.5	“Market mechanics are external, so protocol designers can ignore them.”	219
32	Why These Become Money	220
32.1	As a Store of Value	221
32.2	As Stack	221
32.3	As Telemetry	222
32.4	Trading One Base Reality for Another	222
33	Conclusion: A Bell Labs for Privacy, Proofs, and Compute	223
33.1	The Quiet Mic Drop	223
33.2	Final Thoughts	224
33.3	Equip, Don’t Manage	224
33.4	Two Loops, Not One	225
	Sources	226
A	Formal Model of Verification Asymmetry & VerifyPrice	232
A.1	Definition 1: VerifyPrice(W) and Verification Overhead	232
A.1.1	Implementation Note: VerifyPrice Telemetry Schema (Sketch)	233
A.2	Definition 2: Facility Energy Receipt (FER)	234
A.3	Definition 3: PIDL Receipt (Proof Interface Definition Language)	234
A.4	Definition 4: VerifyReach(N, R)	235
A.5	Definition 5: VerifySettle(C)	236
A.6	Definition 6: Repression Wedge	236
A.7	Definition 7: Liquidation Effect	236
A.8	Definition 8: Physical VerifyPrice	237
A.9	Definition 9: VerifyFlow	237
B	Practical KPIs & Telemetry Templates	238

C The SDK (Proofs-as-a-Library)	240
D Energy & Plant Architecture	243
E Hardware Profiles	245
E.1 Concept	245
E.2 Naming	245
E.3 Hardware Profile Specification	246
E.4 Hardware Profiles in Receipts and Policies	246
E.5 Example Profiles	246
F Communications Resilience Mechanisms	248
G Glossary of Terms & Notation	249
H Formal Model of Market Realization, Wrapper Flows, and Price Capture	252
H.1 Flow Elasticity	252
H.2 Leveraged Rebalancing and Flow-Induced Exposure	253
H.3 Net Mechanical Gain	253
H.4 The Recycling Boundary	254
H.5 Wrapper Recycling Ratio	254
H.6 Return-Decoupled Allocation Flow	254
H.7 Market Impact and Liquidity	255
H.8 Volatility Drag on Leveraged Wrappers	255
H.9 Dealer Hedge Decomposition	256
H.10 Measurement Contract	256
H.11 Reference Stress-Harness Sketch	257
H.12 Boundaries and Caveats	257

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: Communi- cations & 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 Distribu- tion & 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 as AI Money</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 “Private Money” and “AI/Proof Money” actually settle.</i>	Monetary: Chapter 8 <i>Privacy as Private Money</i> , Chapter 11 (Work Credits); Stack: Chapter 21 <i>Layer 5</i> (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)
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); nine public boards (Section 24.9); 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)

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 3.4), **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 33.3).

v1.3 additions: A thread that runs *around* the stack rather than through it. The **Market Realization Plane** (Section 11.13) 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 28.6), and the formal flow model in Appendix H. Readers who want only the conceptual payload can read Section 11.9.1 and Section 33.4 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 15.9 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 31.3.5).

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.

Thesis in Plain Language

The core claim: Three cryptographic capacities—**Privacy** (censorship-resistant settlement that preserves agency), **Proofs** (portable attestations), and **Compute** (verified useful work)—can function as monetary primitives for a dense digital civilization **only if** they are supplied on rails where verification remains cheap, access remains neutral, settlement remains private and refund-safe, and recurring demand for those capacities accrues to a scarce asset through fees, burns, collateral, and disciplined issuance—without the stack itself becoming a surveillance or dependency layer.

The mechanism: Value accrues to holders of triad assets (Work Credits, network tokens) 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 must hold:

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 but not a store-of-value asset.

Demand for these capacities is structural AI needs verified compute, commerce needs private settlement, trust needs proofs so fee revenue persists through cycles.

The falsifiable test: If VERIFYPRICE (verification cost) stays low, VERIFYREACH (reachability under censorship) stays high, and VERIFYSETTLE (settlement success and refund safety) remains robust, the thesis holds. If any of these degrade, the asset becomes just another platform IOU.

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:** the triad (Privacy, Proofs, Compute) as a next-generation store of value (Private Money & AI Money).
- **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 11.13](#)).

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)

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.

Redemption semantics: Implementation-dependent. Work Credits can be designed across a spectrum:

- **Non-redeemable but scarce:** pure SoV instruments where credits represent historical work (like BTC tied to historical hashes). Value derives from scarcity and demand, not redemption rights.
- **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. Unlike SHA-256 PoW, the work is *useful*. 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.

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: Monetary Object Hierarchy

The thesis distinguishes three 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.

Capacity Objects (limited supply, not primary SoV):

- **WC-Voucher:** Prepaid access to future triad capacity at SLA-guaranteed rates. May expire. Useful for hedging, not for long-term savings.

Monetary Objects (scarce, SoV candidates):

- **Base Token / WC-Base:** The native unit required for fees, staking, collateral, and governance. **This is the thesis's primary store-of-value candidate.**
- **LP/Staking Shares:** Derivative positions that grant fee exposure. Secondary SoV instruments.

Hierarchy rule: Throughout this thesis, "the asset" refers to the **base token** unless otherwise specified. Work Receipts are evidence; WC-Vouchers are capacity hedges; the base token is the monetary object whose store-of-value properties are under examination.

v1.4 Quick Reference: Physical Sovereignty & Substrate

The v1.4 update makes the physical layer falsifiable and names 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 15.9.2): probability-weighted verification cost across physical disruption states; resilience is subordinate to constitutional VerifyPrice SLOs.
- **Red Line 13: Energy Sovereignty Failure** (Section 28.5): 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.

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.

v1.2 Quick Reference: Political-Economy and Physical-Capacity Definitions

The v1.2 update extends 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 3.4): 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 3.4, 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 11.12): the ratio of custodial or synthetic economic exposure to protocol-native usage. A rising WDR indicates the asset may be financializing faster than it is becoming money.

v1.3 Quick Reference: Market-Realization Definitions

The v1.3 update adds a 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 7.2): native instruments (Work Credits, staking, LP, corridor claims) touch the protocol and exercise fees, burns, or collateral. External financial wrappers (spot ETFs, ETPs, treasury companies, custodial balances, futures, options, swaps, leveraged ETPs) do not. 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 11.13): 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): $L(L - 1) + \epsilon L^2$, the coefficient converting an underlying return into a wrapper's return-coupled exposure demand.
- **Wrapper Recycling Ratio (WRR)** (Appendix H.5): $-\epsilon L / (L - 1)$; equal to 1 when holder flows exactly offset gross rebalancing, below 1 when the wrapper amplifies, above 1 when it is countercyclical.
- **Return-coupled vs. return-decoupled flow** (Appendix H.6): return-coupled flow responds to the return and shapes volatility and persistence; return-decoupled flow arrives independent of it and shapes destination and level.
- **Wrapper-Native Growth Gap (WNG)** (Section 24.8): $\Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$; 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.

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 “cyberpunk 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 seven 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 is money), [Section 28.5](#) (the thirteen 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, [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 was a brilliant bootstrap: it minted digital scarcity by tying consensus to thermodynamic cost. The next stage in the evolution of blockchains keeps what made PoW legitimate (open admission and unpredictable leader election), but swaps the work function so the “lottery ticket” is earned by producing succinct, publicly verifiable receipts of useful compute. Miners don’t win by burning cycles on nonce search; they win by attaching a proof that some market-demanded computation was done correctly.

A natural anchor workload is matrix multiplication. MatMul-PoW 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, not heat no one can use.

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 stop subsidizing waste and start underwriting capacity the world already buys: ZK proving, matrix multiplication, verified inference. That is the philosophical and operational upgrade: a puzzle that mints receipts, not heat, so **Privacy, Proofs, and Compute may begin to earn monetary premium**—if the conditions outlined in this thesis hold.

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 (energy, hardware, jurisdiction) become un-

verifiable; (9) AI compute enclosure by a small number of hyperscalers persists; (10) lawful privacy collapses socially through systemic forced disclosure; or (11) verification affordability becomes a sovereign policy variable through energy rationing, curtailment, or hardware denial. These are illustrative; the full set of thirteen 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 can function as monetary primitives.**

Privacy (censorship-resistant settlement that preserves agency), **Proofs** (portable attestations of computation and provenance), and **Compute** (useful work wrapped in succinct guarantees) are scarce capacities the world must continually buy. Each can be engineered to be credibly scarce, permissionless, censorship-resistant, and cheap to verify.

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 next reserve asset is not a single object but a **triad of verifiable necessities**: Privacy, Proofs, and Compute. Treat this triad explicitly as cypherpunk monetary primitives:

- **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 assets: what cannot be forged, what no one has to bless, and what everyone can check. They are *scarce capacities* the world must continually buy because life in a dense, digital civilization requires them. Each has its own economy; together they form a monetary base.

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. Where gold condensed geology and Bitcoin condensed randomness, 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.

Networks that supply these at scale will earn a durable store-of-value premium because the world must keep buying their utility through every cycle, even (and especially) under financial repression and currency debasement.

The thesis will establish this through a conditional chain: utility demand creates standardized work; standardized work produces receipts; receipts enable markets; markets produce fee flows; fee flows support a scarce asset *only under non-bypassable accrual rules*; and if that asset remains liquid, neutral, verifiable, and legally holdable, it may 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. 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 31.1.2).
- P5. 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 28.5).
- P6. The system must remain falsifiable.** Verification cost, reachability, settlement safety, decentralization, fee coverage, native-asset capture, physical-substrate legibility, and agency must be public telemetry—not marketing claims (Chapter 24, Section 28.5).
- P7. 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, or VerifySettle.

Conclusion (conditional). Privacy, Proofs, and Compute can support next-generation store-of-value instruments *if and only if* the seven-layer stack converts indispensable digital utility into scarce, verifiable, non-bypassable, agency-preserving asset demand—and if public telemetry shows both the protocol loop and (where relevant) the market-representation loop remain healthy.

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 “cypherpunk 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 Privacy, Proofs, and Compute as monetary primitives and for defining “AI money” and “ZK money” as assets backed by verifiable work.

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).
- 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.

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.

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.

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 monetary primitives. The task is not to find a new story to believe in, but to build a stack where cheap, public verification and indispensable utility do the work that soft guarantees used to perform.

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

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 $\sim 100\%$ of GDP by decade-end, and interest costs compounding, repression is arithmetically attractive to policymakers versus explicit default or politically costly aus-

terity.

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 11.12 formalizes this substitution risk as the **Wrapper Dominance Ratio**.

3.2 The Repression Playbook, Then and Now




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 2½% on long bonds; dashed line marks the 1951 Treasury–Fed Accord. Source: Rose (2021), Federal Reserve Bank of Chicago.

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




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.

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.

3.3 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 (e.g., $O(n^2)$ verify vs. $O(n^3)$ produce; " $\sim 1 + o(1)$ ")

denotes a constant-factor overhead relative to the best known verification baseline). 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.

3.4 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.

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 33.3](#) 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.
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** issuance schedules and realized issuance vs. schedule for triad assets and Work Credits.
- **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 11.12).
- **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 fee+burn share of the security budget remains $\geq$ a target threshold (e.g., 40–60%) and **VerifyPrice** (p95) remains $< 5s$ for core workloads.
- **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 collapse 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**.

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 $p95 \leq 5s$, $r(W) \leq 0.3$	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

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 28.5](#) 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 can function as monetary primitives when verification is cheap and public. A threat model makes that claim 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 IXP's into enforcement arms.

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

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 11.12](#).

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 11.13](#)) exists to make visible, and that VerifyFlow ([Section 24.8](#)) instruments.

With this addition the threat model now spans three distinct kinds of failure: **hostile capture** (state and platform coercion), **institutional co-option** (administrative repression and wrapper substitution), and **endogenous mechanical fragility** (compositional adversaries). A design that defends against only the first is defending against the least likely.

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:

Playbook Move	What It Breaks	Counter
DNS/IP/SNI/QUIC filtering & active probing	Reaching verifiers/bridges; prove-nance 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.

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** triad assets with no fixed coupons, Work Credits priced by market demand for Privacy/Proofs/Compute rather than by decree.
- **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.

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:

- **Curtailement 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 15.9.1](#) 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 28.5](#)) 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:

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

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.

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 3.4 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.

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.

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 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 we want Privacy, Proofs, and Compute to behave like monetary primitives rather than platform features, we have to meet them 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.

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 poli-

tics 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 nervous 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 com-

modities 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 as Monetary Base

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 store-of-value instruments if verification remains cheap, privacy settlement remains usable, useful-work markets avoid capture, and protocol design converts recurring triad demand into scarce, non-bypassable asset value.** They are not simply features of a protocol; they are verifiable necessities that economic actors must keep buying every cycle. This Part explains how that triad behaves as a store of value (SoV), and how **Work Credits** turn verifiable work into energy-anchored monetary instruments (Private Money and AI Money).

Chapter 7

The Triad as Monetary Base

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 **monetary base** in this frame is not “the token” but the **networked capacity to deliver Privacy, Proofs, and Compute under adversarial conditions**. Tokens, credits, and instruments are just ways of slicing claims on that capacity.

A few distinctions help:

Base vs. wrappers: The *base* is “how much verifiable Privacy/Proofs/Compute per unit time this stack can deliver at target SLOs.” *Wrappers* are things like Work Credits, corridor tokens, and staking positions that slice this capacity into transferable claims.

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?”

From a monetary perspective, the key property is that **the world must keep buying the triad’s utility**:

- 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 *replaces* all money, but that it behaves like a **reserve asset**: a base layer of verifiable capacity that other instruments reference, hedge with, or settle into. Private Money and AI Money are just two lenses on how that base is held and used.

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 → markets → fee flows → a scarce asset (only under non-bypassable accrual) → 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 28.5).
4. **Receipts enable markets.** Cheaply checkable receipts 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.
5. **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.
6. **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 11.10, Section 11.11).
Severed if: users obtain equivalent triad capacity through hyperscalers, stablecoins, or custodial APIs without the asset—Red Line 6 (Section 28.5, Section 26.2).
7. **A scarce, captured 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).
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 seven links would establish that the asset *deserves* monetary premium. 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 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 5 all consume bulk power, and link 3—the hinge—is denominated in real resources on reference hardware. That makes the chain dependent on conditions that states partly set: energy availability and pricing, and the obtainability of unprivileged reference hardware (Section 5.5, Section 15.9.1). The substrate exposure is aggregated as sovereign optionality (Section 15.8) and made falsifiable by Red Line 13. The chain is not suspended in mathematics; it is bolted to a grid.

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

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
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
Non-bypassability	Strong	Explicit red line and operator checklist
SoV premium	Strong (conditional)	Nine requirements, thirteen red lines

Table 7.1: Link-by-link strength of the conditional chain. The two medium links are engineering and market-structure risk, not logical gaps.

The rest of this Part defends links 6 and 7, which are where the monetary argument actually lives. Parts III and IV defend links 2 through 5 by building the stack that supplies them. Parts V and VI make every link falsifiable.

7.2 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:

Instrument	What It Represents	Scarce?	Transfer?	SoV?
Base Token	Native unit; required for fees, staking, governance	Yes	Yes	Primary
Work Credits	Claims on verified capacity; minted against work	Yes	Yes	Secondary
LP/Staking	Rights to fee share from corridors, pools, validators	Yes	Sometimes	Derivative
Capacity Vouchers	Prepaid access at fixed rates	No	Yes	No
PIDL Receipts	Proof that work was done	No	Yes	No

Table 7.2: Triad instruments and their SoV candidacy.

The **primary SoV candidate** is the **base token**: the unit in which fees are paid, burns occur, and collateral is posted.

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: Work Credits, staking and LP positions, corridor claims, capacity vouchers. They touch the protocol. Using them exercises fees, burns, or collateral.
- **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.

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.

Object layer	Example	Touches protocol?	Path dependence	Monetary treatment
Base capacity	Privacy, proof, compute through-put	Yes	No	Underlying monetary base
Native monetary object	Base token / WC-Base	Yes	Low	Primary 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 7.3: 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.8).

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): Locks supply proportional to network size.

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:** New tokens cannot be minted arbitrarily; issuance is tied to capacity growth.

Net effect: At steady state, fee burns $\geq$ new issuance, so total supply is flat or declining. Holders benefit from both yield (if staked) and appreciation (via burns).

7.3 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.
- **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? The thesis is compatible with all three, but **Design A (Base Token as SoV)** is the primary reference design because:

1. It is closest to the BTC model that has demonstrated SoV behavior.
2. Fee burns create clear, measurable scarcity.
3. Staking and collateral create structural demand beyond speculation.
4. VerifyPrice and telemetry make the value linkage auditable.

Design B is useful for enterprises but is explicitly *not* pitched as a SoV. Design C is viable but adds complexity (index construction, rebalancing).

The rest of Part II assumes Design A as the default when discussing “the asset” or “Work Credits as SoV.” Where Design B semantics apply (vouchers, hedging), we will note it explicitly.

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.**

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, Work Credits) can be held as a store 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.

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 instruments:** The “money-like” object is not “a proof” but an instrument whose integrity is enforced by proofs—e.g., Work Credits tied to verified proof workloads.

“Attestation Money” in this thesis refers primarily to **proof-backed instruments and claims on proof capacity**, not to proofs themselves.

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 **Attestation Money**:

- 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.

These instruments function as **Attestation Money**:

“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.

Attestation Money is what you hold when you believe the long run belongs to systems where **verification, not vibe, mediates trust**.

Chapter 10

Compute as AI Money

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.

But raw compute is **not yet money-like**:

- 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 fully support monetary-grade issuance. Probabilistic or TEE-backed inference verification (Tiers B and C) receives haircuts or service-only status. The “AI Money” claim weakens if verified compute trades at no durable premium to ordinary cloud compute. See [Chapter 22](#) for the tier taxonomy.

The triad reframes compute as **AI Money** 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 monetary primitive:

“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’ be a store of 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”: AI Money 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 the verification premium will persist because (a) demand for trustworthy AI is structural, and (b) the compliance/sovereignty premium will grow as AI becomes more consequential. But if telemetry shows otherwise, the “compute as SoV” leg weakens relative to privacy and proofs.

AI Money instruments can include:

- **Work Credits:** minted against verified-compute contributions, redeemable for future workloads or tradeable as a savings asset.

- **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 is this money-like rather than just “another utility token”? 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 unit of AI Money corresponds to a valid proof.
- **It’s duration-neutral:** No fixed coupon; value is backed by the ability to sell compute into whatever nominal budgets exist.

Chapter 11

Work Credits: Energy-Anchored Claims

The triad gives us three capacities. To turn them into concrete monetary instruments, we need a **unit of account for work**: something that binds energy, hardware, and verification into a transferable 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,” “AI Money,” and “proof money” are generic, notional instruments (a class of designs that any future chain or protocol could implement), not product branding for a particular network. The point is to specify what *kind* of asset can sit on top of the triad, not to name one specific chain or ticker.

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).

Variant	Rights Conveyed	Scarce?	SoV?
WC-Base (monetary)	No direct redemption. Represents historical verified work. Required for fees, staking, governance. Supply capped; burns create scarcity.	Yes	Yes (primary)
WC-Voucher (prepaid)	Redeemable for future proofs/compute/settlement at SLA-guaranteed rates. May expire or decay.	No (supply tracks capacity)	No (hedge instrument)

Two design variants for Work Credits: Recommendation: The thesis adopts **WC-Base** as the primary reference design for SoV analysis. WC-Voucher is useful for enterprises hedging cost volatility, but is explicitly *not* pitched as a store of value.

Issuance mechanics for WC-Base: 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 does *not* represent a “future option” on its own—that is the WC-Voucher variant. WC-Base represents **historical verified work** and derives its value from being the fee/staking medium, not from redemption rights.

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.
- Unlike SHA-256 PoW, the work is **useful**: 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 something the world actually cares about (anonymized payrolls, authentic media, verified inference), not just burned electricity.

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 SoV: If 90% of Work Credits are minted on L0-A hardware and a major TEE is compromised, the SoV thesis takes a hit—not because the cryptography failed, but because the base assumption (honest machines) was violated.

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 §14.3.1 (Layer 0 Feasibility Ladder). The key point for Part II: **hardware trust is not binary; it is graded, measured, and priced.**

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 is mediated through Work Credits and receipts. The robots and services are paid in Work Credits; in exchange they produce receipts that can be verified cheaply and settled over privacy rails. The more of the economy runs this way, the more natural it becomes to hold savings directly in the unit that pays for these things instead of in a

separate, compliance-backed currency that must be constantly converted.

This is one concrete way to read remarks like “money will just be energy” in a non-hand-wavey way. Energy, compute, and hardware are the inputs that enable robots and AI to work; Work Credits are the units in which their work is paid for and stored. The credits do not magically become Joules; they become a convenient way to denominate future access to verifiable, robot-mediated output.

11.4 Why This Is Still “Money,” Not Just Coupons

If we revisit the SoV requirements from [Chapter 4](#), the mapping is straightforward:

- **Credible scarcity.** Issuance is programmed and publicly documented, either as an explicit schedule (halvings, caps) or as a capacity-linked minting rule. Deviations are detectable on-chain and in telemetry.
- **Cheap, public verification.** Every unit of capacity that credits buy is backed by receipts that any honest party can verify under published VerifyPrice targets.
- **Censorship-resistance & portability.** Settlement uses non-custodial privacy rails and atomic swaps; Layer 1 and 2 ensure clients can reach the network under filtering; Layer 3 ensures users are not doxxed by default.
- **Neutrality & permissionlessness.** Open admission for provers, miners, routers, and mirrors is enforced by decentralization telemetry and fairness tests.
- **Native demand.** The demand for Work Credits comes from real workloads: proofs, verified FLOPs, private settlement, robotics and AI tasks. These are line items in budgets, not speculative self-reference.
- **Lawful privacy by design.** Viewing keys, receipts, and provenance proofs allow institutions to satisfy audits without re-introducing custodians or surveillance chokepoints.
- **Duration-neutrality.** Work Credits do not promise fixed nominal coupons that can be pinned negative by policy. Their value comes from the intersection of energy/hardware costs and buyer budgets for capacity.
- **Asset-level value capture (non-bypassability).** Credits are minted only against accepted work receipts, and the base asset is required for fees, staking, and collateral in the same loop—so triad demand cannot be served at equal quality through a channel that never touches the monetary object. This is the requirement most likely to fail in practice, which is why it is stated as [Value Capture Lemma \(Lemma 11.1\)](#), enforced in [Section 11.11](#), and retired by Red Line 6 rather than left to good intentions.
- **Agency preservation.** Credits are bearer claims on capacity, not accounts: they can be held and spent non-custodially, disclosed selectively through PIDL predicates rather than by opening a biography, and used without a global identity ([Item 9](#)). A credit system that could only be used through custodial intermediaries and mandatory disclosure would sat-

isfy every requirement above and still fail this one.

Two of these deserve emphasis because they are the ones a merely useful capacity token would fail. The first seven describe a well-designed commodity claim; the last two are what make it a candidate store of value rather than infrastructure exposure.

Seen in this light, Useful Work Money is not a loyalty program or a scrip. It is a monetary instrument whose backing is precisely the capacities this thesis has been concerned with all along: privacy, proofs, and compute, grounded in verifiable machines and energy, and measured by VerifyPrice and FERs. It earns a store-of-value premium not because it tells a compelling story, but because in a dense digital civilization it is one of the few things the world must keep buying, on rails anyone can audit.

11.5 Monetary Role

Work Credits sit at the junction between **base capacity** and **financial representation**:

- For operators, they are **revenue in kind**: miners/provers/routers earn credits by contributing triad capacity.
- For users, they can be **pre-paid capacity** or **savings**: hold credits to ensure future access to triad services, or to speculate on increased demand.

Their SoV behavior 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.
- **Repression intensity**: As repression rises, demand (and thus value) for Private Money credits (privacy/settlement workloads) and Attestation Money credits (provenance workloads) should increase.

In portfolio terms, Work Credits are the **equity-like layer of the triad**; they express exposure to the growth and usage of Privacy, Proofs, and Compute under verifiable rules.

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

A critical question for any SoV claim is: **why does demand not simply expand supply, neutralizing scarcity?**

If Work Credits are minted whenever verified work is done, and demand for triad services grows, supply grows too. This could make WC behave like a **labor-backed scrip** rather than a scarce reserve asset.

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

$$\Delta \text{Supply} = \text{Issuance} - \text{Burns} - \text{Lost} \pm \text{Governance}$$

Term	Definition	Typical Magnitude
Issuance	New WC minted against Work Receipts, subject to issuance envelope	Bounded by schedule or capacity ceiling
Burns	WC permanently destroyed when used for fees	30–50% of fee volume (reference design)
Lost	WC in lost/forgotten wallets	~1–2% of supply per year (empirical from BTC)
Governance	Adjustments via protocol upgrades (rare, requires supermajority)	Near zero in steady state

For WC-Base to behave as SoV, the design must ensure:

$$\text{Burns} + \text{Lost} \geq \text{Issuance} \quad (\text{at steady state})$$

This means:

- **During growth phase:** Issuance $>$ Burns, but capped by schedule/capacity. Supply grows, but at a predictable, declining rate (like BTC halvings).
- **At maturity:** Burns $\geq$ Issuance. Supply is flat or declining. Holders benefit from scarcity.

What prevents runaway issuance?

1. **Issuance envelope:** Total WC mintable per epoch is capped, regardless of work submitted. Excess work earns priority, not extra tokens.
2. **Capacity linkage:** Issuance envelope expands only when new verified capacity (FERs, hardware profiles) is registered. You can't mint by decree.
3. **Halving schedule:** Many designs use BTC-style halvings to ensure long-run issuance approaches zero.
4. **Governance friction:** Changes to issuance rules require supermajority and are visible in dashboards before activation.

Telemetry that detects supply risk:

- **Inflation rate:** Issuance / circulating supply. Should decline over time.
- **Burn rate:** Burns / fee volume. Should stay within target range.
- **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.

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 SoV, Not Just a Utility Token

It is easy to misread Work Credits as another “utility token” with a story stapled on top. §3 and Section 11.4 already showed that they inherit the full SoV checklist (scarcity, cheap public verification, censorship-resistance, neutrality, native demand, lawful privacy, and duration-neutrality) from the stack. Here we make that contrast explicit.

Four properties do most of the work:

1. Issuance is constrained by energy and installed capacity.

Work Credits do not float on arbitrary governance votes. Issuance is bounded by measured energy and verifiable capacity: FERs and plant telemetry cap the budget of subsidized work over a given interval. To print more, someone has to build more verifiable machines or increase real capacity. Narrative cannot override thermodynamics.

2. Demand is driven by budgeted workloads, not vibes.

The “use” of a Work Credit is paying for proofs, inference, settlement, and related workloads that already sit on OPEX lines. AI labs, exchanges, custodians, and regulators must keep buying these workloads each quarter, regardless of token price. That places Work Credits *upstream* of compulsory spend, not downstream of hype.

3. Verification asymmetry keeps cheating costlier than honesty.

For each workload W , the verification ratio $r(W)$ is explicitly bounded: it is always much cheaper to check a PIDL receipt than to fake it at scale. A Work Credit backed by a portfolio of such workloads inherits this property: to counterfeit the base, you must fake receipts, and faking receipts is provably more expensive than doing the work. Most “utility tokens” rely on soft notions of “activity”; Work Credits rely on receipts that anyone can cheaply reject if they are bogus.

4. Returns are duration-neutral and fee-driven, not coupon-driven.

Work Credits do not promise fixed coupons or redemptions at par. Their economic value comes from a share of fee+burn and spread on a growing, compulsory workload budget. As more proofs, compute, and private settlements clear over the stack, more fees flow in and more units are retired. That makes them duration-neutral: they behave more like equity in a necessary utility than a bond whose real yield can be pinned negative.

Taken together, these properties push Work Credits into the store-of-value bucket described in §3 and Section 11.4:

- constrained by energy and capacity,
- demanded by necessity rather than fashion,
- protected by verification asymmetry, and
- paid in a fee stream tied to indispensable workloads.

In that sense they are not tickets to a theme park; they are metered claims on the **power plant** that keeps the park (and the rest of the AI + ZK economy) online.

Work Credits are “utility tokens” only in the most literal sense: they buy utility the world cannot stop buying. Under the SoV lens, that is exactly what you want: an asset whose backing is quantitatively visible in receipts and KPIs, and whose long-run value is rooted in capacities that must be purchased through every regime.

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: **supply is bounded by physics (energy, hardware) and shrinks with usage (burns), while 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 causal chain:

1. Demand for Privacy/Proofs/Compute
2. ↓ Users must acquire native tokens to pay fees

3. ↓ Fees paid → portion burned, portion to stakers/provers
4. ↓ Burns reduce supply; staking rewards require holding
5. ↓ Increased demand + reduced supply → price appreciation
6. ↓ Holding the asset captures the economics of the triad

More concretely:

- **Fee revenue:** Every proof, every private settlement, every verified inference pays a fee denominated in the native asset. This creates continuous buy pressure.
- **Burn mechanics:** A fraction (e.g., 30–70%) of fees is burned. As throughput grows, more units are destroyed than minted, creating deflationary pressure during growth phases.
- **Required collateral:** Provers, routers, and LPs must stake tokens proportional to their capacity. This locks supply and aligns incentives with network health.
- **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 supply shrinks as those workloads grow. 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:

Risk	Mitigation
Operator rent extraction	Fees are split: burn + stakers + provers. Operators cannot capture 100%; a structural portion goes to asset retirement.
Governance inflation	Issuance is constrained by FERs and capacity telemetry. Minting beyond real capacity triggers SLO breaches visible in dashboards.
Holder dilution	Burns offset new issuance. Net supply is designed to be flat or declining during steady-state usage.
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.

The economic design goal: At steady state, fee burns $\geq$ new issuance, so holders benefit from both yield (staking rewards) and capital appreciation (supply contraction).

Summary: The Value Accrual Lemma

Demand for triad capacity (Privacy, Proofs, Compute) translates into store-of-value premium for the native asset because:

1. **All usage requires the asset** (fee medium).
2. **Usage burns supply** (deflationary pressure).
3. **Capacity provision requires staking** (supply lockup).
4. **Issuance is energy-constrained** (no governance inflation).
5. **Workload demand is structural** (AI, commerce, compliance budgets, not hype cycles).

If these five conditions hold, and VerifyPrice/Reach/Settle remain within SLOs, holding the asset captures the economics of the triad rather than merely consuming a service.

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 not a store-of-value asset. This lemma is the central bridge between “these capacities are indispensable” and “therefore a specific asset earns monetary premium.”

This is the bridge between “these capacities are indispensable” and “therefore they can function as monetary primitives and earn a durable store-of-value premium.”

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

monetary adoption $\nRightarrow$ 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 $\rightarrow$ native fees $\rightarrow$ burns, collateral lockup, and operator demand $\rightarrow$ scarcity and monetary premium. It is governed by [Lemma 11.1](#).

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 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:

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.** Every use of the triad must flow through the base token. If operators accept fiat directly, the asset is optional.
3. **Fees must produce burns, retirement, or staking yield.** Fee revenue must reduce circulating supply or compensate holders, not merely pay operators.
4. **Operators must post the asset as collateral.** Provers, routers, and LPs must lock the asset to participate, creating structural demand beyond fee payment.
5. **Issuance must be capped or capacity-constrained.** New supply cannot expand by governance decree; it must track real, verified capacity.
6. **Users must not be able to bypass the asset at equal service quality.** If AWS, a ZK prover marketplace, or a stablecoin-based privacy wallet can sell equivalent triad capacity for fiat, the native asset becomes unnecessary.
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 condition fails, the asset is infrastructure exposure, not money.** It may still be useful and valuable, but it does not earn a store-of-value premium distinct from its service value.

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 can earn monetary premium. The [Value Capture Lemma](#) formalizes these conditions; the telemetry regime ([Chapter 23](#)) makes them falsifiable.

11.11 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 Work Credits entirely?*”

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, creating structural demand beyond speculation.
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).

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 exposure 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).

11.12 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:

$$\text{WrapperActivityRatio}_{t,h} = \frac{V_{t-h:t}^{\text{wrapper}}}{V_{t-h:t}^{\text{native}}}$$

Both are bounded and 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.

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 28.5](#) and tracked on a dedicated Wrapper Dominance Board ([Section 24.9](#)).

11.13 The Market Realization Plane

The Wrapper Dominance Ratio measures a stock: 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.

11.13.1 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, AI concentration, and compute demand. 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 [Lemma 11.1](#) restated as a coefficient.

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 fundamental or native monetary impulse, Q_t is aggregate exposure demand, 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.

Chapter 12

Monetary Design Space

Once Work Credits exist as a generic mechanism for binding work to claims, we can sketch the **design space of monetary instruments** built on top of the triad.

12.1 ZK Money

ZK Money instruments primarily reference **Privacy + Proofs**:

- Claims on shielded settlement capacity (corridor LP positions, privacy-rail credits).
- Claims on zero-knowledge proof capacity (e.g., SNARK/STARK service tiers).

They behave like:

- **Private Money**: bearer-like, censorship-resistant settlement rights.
- Plus **Attestation Money**: bundled attestations of compliant behavior.

ZK Money is attractive to:

- Treasuries under repression.
- Regulated actors who need both privacy and attestable compliance.
- Individuals who want “cash-like” assets in digital form.

12.2 AI Money

AI Money instruments primarily reference **Compute + Proofs**:

- Claims on verified FLOPs.
- Futures on model-specific inference capacity with proof guarantees.
- Equity-like positions in proof factories whose output is standardized, verified compute.

They are attractive to:

- AI labs and application builders wanting hedges against compute shortage.
- Investors who believe “AI demand will outlive this particular model cycle.”

12.3 Hybrid Instruments

Hybrids reference all three legs:

- **Triad baskets:** Index-like instruments backed by diversified Work Credits across privacy, proof, and compute workloads.
- **Corridor-linked AI Money:** Instruments that couple private settlement rights with AI inference capacity.
- **Agent treasuries:** On-chain agents that hold combinations of ZK Money and AI Money to fund their own operation.

The key design constraint is always the same: **claims must stay tightly coupled to verifiable work and capacity**, not to governance mood or marketing narrative.

Chapter 13

ZK Money and the ZK + AI 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 underpin **ZK Money** and **AI Money** 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” – AI Money substrate):** Duplex-style and Ambient-style PoW systems treat model training and inference as work functions. Work Credits minted against these workloads become **AI Money**.
2. **ZK blockchains (“proof chains” – ZK Money substrate):** Nockchain-style zk-PoW systems coordinate global prover markets. They mint **ZK Money** (claims on future proof capacity).
3. **Privacy rails as connective tissue (Private Money substrate):** 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 a particular triad-based instrument is a credible store of value.

For any candidate instrument X (a Work Credit type, a ZK Money basket, an AI Money token), ask:

1. **What triad workloads back it?**
 - Exactly which Privacy, Proof, and Compute flows does X reference?
 - Are those workloads canonical and well-specified?
2. **How is issuance tied to work?**
 - Is issuance formulaic and transparent (e.g., one credit per verified unit of W)?
 - Can governance mint outside of those rules? Under what constraints?
3. **Can anyone verify the backing?**
 - Are VerifyPrice and workload telemetry public?
 - Can a third party recompute or sample the proofs/settlements that justify supply?
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?
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 15.9.2)?
- Does the resilience premium demonstrably lower disruption-adjusted cost, or is it capital misallocation with a security narrative attached (Section 15.9.3)?

*Note the distinction from the previous question, and do not let an operator collapse them. The prior question asks whether the substrate can be *seen*; this one asks whether what we see can *withstand* anything. A facility can be perfectly transparent about being catastrophically fragile and pass the legibility test completely. That gap is why 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 11.12)?
- 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 triad-based instrument that scores well on these questions behaves like a **next-generation SoV**:

- 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 Triad

For context, it is useful to compare how gold and silver—the canonical “old-world” stores of value—score against the same framework:

Criterion	Gold / Silver	Triad (Work Credits)
What backs it?	Geological scarcity; 5,000-year Lindy effect	Verifiable work: privacy settlement, proofs, verified compute
Issuance tied to work?	Mining (energy + geology)	Proof of workload W at tier T ; telemetry-gated
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)	Fee flows from structural demand; no fixed coupon
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. triad-based instruments.

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. However, once VerifyPrice, VerifyReach, and VerifySettle are operational and healthy, the triad offers superior properties for a dense digital civilization: instant verification, programmability, censorship-resistant settlement, and yield tied to structural demand rather than lending risk.

The thesis does not claim gold is obsolete. It claims that for the specific failure modes of the 21st century—financial repression under surveillance, synthetic media, AI-saturated workflows—the triad is better suited to the task. The bridge is still load-bearing; the destina-

tion is under construction.

The rest of the thesis will shift back to the **stack angle** (Layers 0–6) and the **telemetry angle** (VerifyPrice/Reach/Settle, decentralization metrics, governance). But the monetary picture stays in the background: every design choice is judged by how well it supports the triad as a monetary base and how well Work Credits and derived instruments can function as Private Money and AI Money when the world turns adversarial.

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 a **monetary base**, with Work Credits as energy-anchored claims on that 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.”*

That was barely acceptable for **useless work**. It is not acceptable when the puzzle mints **receipts**, not heat; when proofs and verified compute become monetary primitives.

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." A triad backed by machines we cannot interrogate is not a reserve asset; it is an IOU 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)

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 units per quarter	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.

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 $X\%$ of critical capacity.
- **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?”

Profile Criticality	Definition	Min Sample Rate	Funding Source
Tier 1 (Critical)	Top 5 profiles by WC issuance; >10% of total capacity	≥ 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.

Sample rate model (reference design):

Statistical sufficiency: For a profile with N deployed units, sampling n units provides confidence that:

$$P(\text{detect if } \geq k\% \text{ are compromised}) \geq 1 - (1 - k/100)^n$$

For Tier 1 profiles ($n = 50$, $k = 5\%$), detection probability is $\sim 92\%$. This is not certainty—it is **bounded uncertainty**, documented and priced.

What “Gold tier” requires: A profile achieves Gold tier (L0-B or higher) when:

1. $\geq 80\%$ of deployed units have been covered by sampling programs over the profile’s lifetime.
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.

Gold tier unlocks:

- Full WC eligibility (no issuance cap).
- Zero collateral haircut.
- Eligibility for “pristine collateral” designation in DeFi integrations.

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 part of the monetary base.

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 index of the number of independent physical pathways by which verified work can continue under disruption, net of exposure to external coercion. Computed entirely from Facility Capacity Receipt fields.

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, and one term does not survive it.

At sovereign scale the index normally includes an industrial and logistical throughput term. A proving facility has no meaningful analogue: it does not convert feedstock into manufactured output, and borrowing the term would produce a number with a component nobody can measure. We drop it, and say so, rather than carrying a decorative variable.

15.8.2 Mapping to Existing FCR Fields

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 (entropy over independent fuel sources)
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

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, and let H denote a Herfindahl–Hirschman index computed over jurisdictions and balancing authorities:

$$\mathcal{O}_s^{\text{net}} = \sum_f \kappa_f \mathcal{O}_s(f) - w_H H$$

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.

Relationship to existing metrics. H overlaps deliberately with the dispersion metrics already used for verification monoculture and enclosure risk (Section 28.5, 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 28.5, 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 pathway 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 28.5).

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 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 28.5 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.

15.9.1 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.
- **Curtailment, 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.

15.9.2 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.

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:

$$\text{VerifyPrice}_{\text{adj}}(W) = \sum_s \pi_s \cdot \text{VerifyPrice}_s(W)$$

Disruption-Adjusted VerifyPrice

The probability-weighted cost of verification across physical disruption states, rather than the observed cost under benign conditions.

A resilient fleet has higher $\text{VerifyPrice}_{\text{normal}}$ and lower $\text{VerifyPrice}_{\text{adj}}$ whenever the reduction in disruption-state cost outweighs the benign-state premium. 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 28.5.

15.9.3 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 it lowers disruption-adjusted VerifyPrice, **and** only if benign-state Physical VerifyPrice remains inside the constitutional SLO band.

The second 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.

15.10 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**.

15.10.1 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.

15.10.2 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.

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.

Process guarantees:

1. **Evidence threshold:** S1 requires independent lab confirmation; S2 requires reproducible demonstration; S3 requires active exploitation observed.
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.

15.10.3 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”).

15.11 Stress Tests for Layer 0

Every claim in Layer 0 should map to a testable stress scenario:

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.

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.

15.12 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, 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.

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 $Y\%$ 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 continent	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.

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)	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.

Metrics:

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.

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% triggers investigation
Temporal gaming	Continuous measurements (not snapshots); 24-hour rolling averages published

Table 16.3: VerifyReach adversarial robustness measures.

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 monetary primitives, then the binaries that implement them must keep moving even when networks are hostile.

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 one independent builder reproduces 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 collateral (WC-Base) proportional to the value of credentials they issue. 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 PROOF₂²⁰ 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.** Median and tail verify time and cost per standard identity predicate.

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.

Specifically:

Layer	Minimum Viable Requirement	Falsification Test
L0	At least one L0-B grade profile; sampling for L0-A	If all proving runs on L0-A with <20% coverage → thesis weakened
L1	Encrypted P2P + ≥ 2 fallback transports; VerifyReach with ≥ 3 operators	If >20% ASNs unreachable for >7 days → thesis fails
L2	Threshold-signed releases + transparency log + ≥ 2 builders	If single key compromise can push malicious updates → thesis fails
L3	≥ 3 identity predicates with p95 VerifyPrice <5s; issuer pluralism	If identity requires centralized issuer → thesis weakened

Table 18.3: Minimum viable requirements for Layers 0–3.

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% triggers investigation.

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.

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 ≠ 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**.

Layer 4's goal is simple:

Make $r(W) \ll 1$ for the workloads that matter, and keep it that way in production.

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:

$\text{VerifyPrice}(W) = \{\text{p50 time, p95 time, p50 cost, p95 cost, failure rate, hardware profile mix}\}$

20.3 Canonical Workloads and Proof Types

To avoid an unbounded zoo of bespoke proofs, Layer 4 maintains **canonical workloads**:

- **MatMul**($n, k, m; \varepsilon$) – matrix multiplication at specified dimensions and error bounds.
- **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.

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 with sufficient precision. The following template ensures standardization:

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 $\|C - A \cdot B\|_\infty \leq \varepsilon$.”

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: $O(\log n)$ for succinct proofs; $O(n)$ for non-succinct
- 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	$p95 \leq 60\text{s}$	$1\times$ verification	$1.0\times$
Silver	$p95 \leq 10\text{s}$	$2\times$ verification	$1.5\times$
Gold	$p95 \leq 2\text{s}$	$3\times$ verification + audit	$2.5\times$

Table 20.1: SLA tiers for canonical workloads.

SLA Tiers:

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).

Field	Value
WorkloadID	MATMUL_4096_FP32
Statement	Matrix multiply $C = A \times B$, dimensions 4096×4096 , FP32, $\ C - A \cdot B\ _{\infty} \leq 10^{-5}$
Public Inputs	Hash(A), Hash(B), Hash(C), ε , timestamp
Private Inputs	A, B, C , intermediate products
Verified	Correctness, Bounds, Freshness
Verifier Complexity	$O(\log n)$, $\leq 256\text{MB}$, $\leq 500\text{KB}$ proof
L0 Requirement	L0-A minimum; L0-B+ for Gold tier
VerifyPrice Target	$t_{95} \leq 2\text{s}$ (Laptop-Class)

Table 20.2: Example canonical workload: MATMUL_4096_FP32.

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 $\leq 2\text{s}$ 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 audi-

tors.

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 monetary claims: The “AI Money” and “Compute as SoV” arguments strengthen as these components mature from Experimental to Production. Until Tier A cryptographic verification is practical for large models, inference-backed monetary issuance should be discounted (see Tier taxonomy in [Section 22.1](#)).

Verification Type	Monetary Eligibility
Full cryptographic proof (Tier A)	Full eligibility for WC-Base issuance
Probabilistic audit with explicit error bounds (Tier B)	Discounted or limited eligibility
TEE/vendor attestation only (Tier C)	Service market only; not monetary-grade
Unverified output	No eligibility

Table 20.4: Verification type and monetary 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 now a monetary KPI—it determines whether Work Credits retain their value proposition. That requires a rigorous measurement harness, not just a dashboard slogan.

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: $c = e \times p_{\text{energy}} + t_{95} \times p_{\text{opp}}$, where p_{energy} is the reference energy price (USD 0.10/kWh) and p_{opp} is opportunity cost (USD 0.001/s), both published quarterly.

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_70B	t_{95}	$\leq 30s$	$> 60s$ for 7 days
All	failure rate	$\leq 0.1\%$	$> 1\%$ for 7 days

Table 20.8: Constitutional VerifyPrice targets.

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 Physical VerifyPrice SLOs (Constitutional)

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.

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 15.9.1 separates the verification-side exposure (hardware obtainability) from the proving-side exposure (bulk energy), which behave differently and should not be conflated.

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.

20.7.3 Token-Quoted VerifyPrice (Market Signal, Not Target)

Token-quoted VerifyPrice is a **useful market signal**, but not a constitutional target:

$$\text{VerifyPower}(\text{token}) = \frac{1}{\text{VerifyPrice}(W) \times \text{FeeSchedule}(W)}$$

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 monetary-grade 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:

Tier Rules for Work Credit Issuance:

Fee + Burn + Slashing Logic:

- 70% of fee to prover (reward)

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.

Tier	Verification Type	WC Weight	Collateral Grade
Tier A	Full cryptographic proof (ZK-SNARK/STARK)	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.

- 20% burned (supply reduction)
- 10% to protocol treasury (security budget)

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.
- **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(payload_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.

Corridor Health SLOs:

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
tff(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.

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.

Minimum Viable Corridor Spec: Corridors failing these requirements are flagged as “experimental” and excluded from default PRK routing.

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.

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

The twelve primitives 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.

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.

Inference Proof Tier Taxonomy Verified inference is the most economically significant—and currently most fragile—part of the “AI Money” claim. To prevent readers from concluding “so AI Money is backed by trust,” we tier the guarantees explicitly:

Why this matters: Without this taxonomy, “verified inference” conflates cryptographic

Tier	Verification Type	Description	Error Bound	WC Weight
Tier A	Cryptographic correctness	Full ZK proof of inference (ZKML)	0 (exact)	1.0×
Tier B	Probabilistic soundness	Audited transcripts + randomized checks + peer prediction	Explicit (e.g., 99% confidence)	0.6×
Tier C	Attestation-backed	TEE attestation + sampling audits	Implicit (trust TEE + sampling)	0.3×

Table 22.1: Inference proof tier taxonomy.

Use Case	Minimum Tier	Rationale
WC-Base issuance (full weight)	Tier A	Monetary issuance requires highest assurance
WC-Base issuance (discounted)	Tier B	Acceptable with explicit error bounds and discount
Service market only	Tier C	Can sell inference; cannot mint WC-Base
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.

proof with probabilistic audit with vendor attestation. By making tiers explicit and tying them to WC issuance weights, the thesis acknowledges engineering reality (full ZKML isn't always practical) while preserving the monetary story (only high-assurance verification backs money-like instruments).

Tier B requirements (reference): For Tier B verification to qualify:

- **Transcript determinism:** Given the same model + input, transcript must be reproducible.
- **Audit rate:** $\geq 10\%$ of inferences are independently re-executed by random auditors.
- **Peer prediction:** Multiple independent provers; divergence 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 ”).

Migration path: As ZKML matures: (1) Tier A coverage expands to more model classes; (2) Tier B issuance weight increases toward 1.0× 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↔ZEC/XMR settlement with refund-safety.
8. **Lawful-privacy corridor pack** — Hooks for viewing keys + receipts enabling regulated actors to prove compliance.
9. **Bridge-safety templates** — Pattern book for cross-chain settlement avoiding “magic multi-sigs.”

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.

Between them the four applications stress eleven of the thirteen red lines. The two they do not reach—Red Line 4 (telemetry capture) and Red Line 8 (governance capture)—are protocol-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.

22.3 Reference Application: Private Treasury & Payroll

Scenario. A globally distributed company wants to pay staff and vendors while preserving salary confidentiality, avoiding unnecessary exposure of the transaction graph, and giving

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 28.5) each would trip first.

auditors and tax authorities scoped evidence that payments were authorized, complete, and properly reported.

Flow:

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 $\leftrightarrow$ ZEC/XMR corridors.
5. **Verify.** Auditors verify correctness via proofs and receipts.

22.4 Reference Application: Media Provenance

Scenario. A media network guarantees that “gold channel” content is accompanied by verifiable provenance.

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

Scenario. An AI service offers “verified inference” to enterprises who don’t want to trust a black-box API.

Flow:

1. **Create.** Client submits inference request with constraints.
2. **Compute.** Proof factory runs the inference.
3. **Prove.** PAL compiles into Inference(M, χ ; policy) with proof.
4. **Settle.** Payment via PRK contingent on valid proof.
5. **Verify.** Client verifies proof and accepts output.

AI Money is just Work Credits and derivatives for these inference workloads.

22.6 Reference Application: Proof/Compute Procurement

Scenario. A DAO wants to pre-buy proof and compute capacity for future needs.

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; Work Credits held in treasury.
5. **Verify.** Anyone audits backing via VERIFYPRICE, profiles, and proof metadata.
The DAO's treasury becomes **triad-backed**: part BTC/ETH, part ZK Money, part AI Money.

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.

Only with these in place does it make sense to talk about **Private Money** and **AI Money** as more than metaphors. 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 can function as monetary primitives; that a seven-layer cypherpunk stack can supply these primitives; and that Layers 4–5 turn work into truth and truth into private, non-custodial settlement.

Part V looks at the same system from the **operations and governance** angle. If the triad is to behave like a store of value under repression, then **“we promise” is not enough**. Neutrality, censorship-resistance, and verification economics must be **measured** and **governed** in ways that are falsifiable, not vibes.

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

The second job of Layer 6 is to define the **control surfaces**: the knobs that can be turned, by whom, and with what receipts.

Economic Parameters:

- Block reward curves and Work Credit issuance schedules.
- Fee policies (floor, burn, split between security and dividends).
- PoUW weighting across workloads.

Technical Parameters:

- Circuit versions and proof-system parameters.
- Accepted hardware profiles (additions, deprecations).
- Corridor policies (timelocks, refund windows, anonymity thresholds).

Operational Parameters:

- Telemetry granularity and reporting requirements.
- Incident severity levels and response playbooks.
- Key ceremonies and quorum sizes.

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 28.5](#)). 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, not waste. The security budget is tied to canonical workloads with published VERIFYPRICE , not to arbitrary hash puzzles with no external demand.

Inference Verification Tiers and WC Eligibility: To prevent the “AI Money is backed by spot checks” critique, inference 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 “soundness haircut” reflecting error bounds and audit rate. 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**:

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.

Requirements: ≥ 3 of 5 security signers; auto-expire in 7 days; postmortem within 14 days.

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 $\rightarrow$ high fees $\rightarrow$ high burn $\rightarrow$ supply reduction $\rightarrow$ value appreciation. This closes the loop between utility and asset value.

Cross-reference: This fee routing supports the [Value Capture Lemma](#) from [Chapter 7](#). 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, WC-Base is retired in specific circumstances:

23.6.4 Risk Haircuts

Work Credits minted under weaker assurance levels receive discounted issuance:

These haircuts are protocol parameters (not discretionary); they appear in dashboards and apply automatically.

Trigger	Retirement
WC-Voucher redemption	Underlying WC-Base used to fulfill capacity claim is retired
Slashing	100% of slashed collateral is burned (not redistributed)
Failed workloads	Fees for failed proofs (prover fault) are burned, not paid

Table 23.5: WC-Base retirement triggers.

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 15.9.1: 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 15.9.3, 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:

$$\text{FeeCoverage} = \frac{\text{FeeRevenue} + \text{Burn}}{\text{TotalSecurityBudget}}$$


Why this matters: A system that relies entirely on issuance is inflating its way to security. A system where fees cover the budget has structural demand. The trajectory from issuance-funded to fee-funded is the path from “speculative token” to “productive asset.”

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

Target	Timeline
FeeCoverage $\geq$ 30%	Year 1
FeeCoverage $\geq$ 50%	Year 3
FeeCoverage $\geq$ 80%	Year 5+

Table 23.7: Fee coverage trajectory.

23.6.6 Value Capture Loop



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

“No dashboards, no trust” is the mantra. But dashboards themselves can be gamed. If Layer 6 claims “measurement is the constitution,” then the measurement regime itself must be bullet-proof.

24.1 Reference Verifier Classes

All Verify* metrics are 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 p95 $\leq$ 5s”) are defined against **Laptop-Class** unless otherwise specified. This ensures the “anyone can verify” promise is falsifiable.

24.2 Reference Verifier Implementations

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.

Why this matters: A single verifier implementation can have bugs that inflate or deflate measurements. Two independent implementations provide a cross-check.

24.3 Cost Vector Definition

VERIFYPRICE is not a single number. 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 selective reporting: only show receipts from good regions, fast hardware, or favorable conditions. The sampling regime prevents this:

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 Multi-operator	Rolling 24-hour windows; no periodic snapshots that can be gamed ≥ 3 independent measurement operators; divergence $> 5\%$ triggers investigation

Table 24.3: Sampling mechanisms for Verify* telemetry.

Rule: Any dashboard point published without a verifiable sampling seed and stratification breakdown is not part of the official Verify* record.

24.5 Data Availability and Anti-Memory-Hole

Dashboards are views. The source of truth is the **receipt corpus**. Receipts must be:

Why this matters: If a government or CDN can delete “bad months,” the telemetry regime is captured. Content-addressing + neutral anchoring makes deletion detectable.

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.

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 11.13 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; 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

$$\text{VerifyFlow}(X) = (Q^{RC}, Q^{RD}, \text{MPR}, \text{CCR}, \text{SER}, \text{WNG}, \text{RCR}, \text{NUS})$$

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, $Q^{RD} = \sum_k w_{Xk} F_k$, where F_k is net creation into wrapper k and w_{Xk} its exposure weight to X . This term shapes destination, concentration, and level — which prices get accepted without valuation-sensitive selling.

Mechanical Pressure Ratio (MPR). Let $Q^{\text{exp}} = Q^{RC} + Q^{RD}$ be total product-level *exposure demand*. Dealers translate that exposure into market orders, so realized market demand is its intermediated image, $Q^{\text{mkt}} = \mathcal{H}(Q^{\text{exp}}, \text{dealer inventory, internalization, swap/future/option mix})$, and $\text{MPR} = |Q^{\text{mkt}}|/V_X$ measures flow pressure relative to executable liquidity V_X . 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.

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). $\text{WNG} = \Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$, where U^{native} counts

native fee payments, burns, collateral, private settlement, proof purchases, and verified-compute consumption. A persistently positive WNG means financial exposure is growing faster than monetary use — the single most compact measure of financialization outrunning adoption.

Recursive Claim Ratio (RCR). Exposure held in wrappers whose underlying is itself another wrapper or derivative.

Native Use Share (NUS). The share of observable economic activity that actually touches the protocol-native monetary loop.

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. 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 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.

24.9 The Public Boards

Metrics surface as public “boards.” The original five have grown to nine as the thesis has extended from pure verification economics into physical capacity, political economy, 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). 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.

Wrapper Dominance Board. Tracks whether the asset is becoming a financial wrapper rather than money: the **Wrapper Dominance Ratio** (Section 11.12) 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, and holder-cohort statistics.

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 the native asset becoming money?*
- **Market Realization & Wrapper Board:** *What is currently setting its price?*

The Value Capture Board stays focused on protocol economics. VerifyFlow stays focused on market realization. A reader who can see both at once can tell adoption from financialization; a reader who sees only one cannot.

These boards are not marketing pages; they are part of the protocol's public interface.

24.10 Detecting Drift

Telemetry detects **drift**:

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.

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 remain usable enough that Work Credits still behave as stores of value.”

Telemetry makes these hypotheses testable:

- If VerifyReach shows certain flows consistently blocked and no alternatives emerge, neutrality 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

Layer 6 sits where protocol meets law and politics. 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.

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	Cannot Require
AML screening	ZK proof of “sender $\in$ cleared-set”	Full sender identity
Jurisdictional limits	ZK range proof of amount	Exact amount
Tax audit	Time-bounded viewing key	Universal history
Counterparty verification	ZK set-membership proof	Real-name disclosure

Table 25.1: Lawful privacy: what can and cannot be required.

25.2 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.”

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.

25.3 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.

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 11.12) 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 monetary base, built a seven-layer stack to supply it, and wired in governance and telemetry so neutrality and repression-resilience are falsifiable. Part VI asks the unromantic question:

What happens when this thing hits the real world? Who uses it first, how does it fail, and under what conditions does it quietly become “money”?

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, and VerifySettle are live and healthy, the triad offers superior properties for a dense digital civilization:

- **Verifiability:** Gold’s purity requires assay; proofs verify in milliseconds on a laptop.
- **Programmability:** Gold cannot enforce conditional logic; Work Credits can encode SLAs, escrow, and policy.
- **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.
- **Yield on utility:** Gold earns nothing unless lent (introducing counterparty risk); triad assets earn fees from structural demand.

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 Four Phases of Adoption

The stack is not adopted by “the market” in one step. It’s pulled in by different constituencies, each with their own pain.

We sketch the curve as four overlapping phases. 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.

27.2.1 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+burn covers $\geq 30\%$ of security budget 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:

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.

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 $\rightarrow$ I regression

Table 27.2: Failure gates for adoption phases.

27.2.2 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.

27.2.3 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 $\rightarrow$ Prove $\rightarrow$ Settle $\rightarrow$ 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) show that verification overhead can remain $r(W) = v/p \leq 0.3$ on commodity hardware and that miners without favoured hardware profiles can still win blocks.
- **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 $\text{VERIFYPRICE}(W)$ dashboards for $W \in \{\text{PROOF_2}^{20}, \text{MATMUL_4096}, \text{INFER_LM_70B}\}$ with months of history, and p95 times stay within stated SLOs under stress.
- **VERIFYSETTLE is measured for at least one serious corridor.** For some BTC↔XMR or BTC↔ZEC corridor C , $\text{VERIFYSETTLE}(C)$ 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?”

27.2.4 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↔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+burn covers a meaningful slice of security budget.** For at least one serious PoUW/proof network, fees and burns tied to real workloads cover $\geq 30\%$ of miner/prover revenue over 12–24 months.
- **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 institutional 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.

27.2.5 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 “flippening.” 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, begin to trade at a **durable store-of-value premium** alongside Bitcoin.

Phase III telemetry and triggers: Phase III is real once:

- **Non-trivial collateral and reserve usage.** Work Credits appear as collateral and reserves in ≥ 3 independent venues: lending markets, derivatives, structured products, custody mandates. A visible share of supply is locked in long-duration roles.
- **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 holding behavior.** Holder data shows a meaningful cohort who hold primarily for access to future capacity and SoV properties, with holding periods measured in years.
- **Fee coverage $\geq 50\%$.** Sustained over multiple quarters.

When these show up in the dashboards and market structure, the triad has joined the shortlist of things serious allocators treat as monetary base, not just tech.

27.2.6 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.

The triad earns a store-of-value premium 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.

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.

27.2.7 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.

27.3 Failure Gates

Phases can regress. If Phase II metrics collapse, the system slides back to Phase I.

Condition	Consequence
VERIFYPRICE p95 exceeds SLO for ≥ 3 months	Phase regression; SoV thesis weakens
Corridor refund_safe < 100% (repeated)	Phase regression; Private Money claim fails
Verification concentrates > 70% on single profile	Decentralization claim fails
Receipt data becomes unverifiable	Telemetry capture; thesis unverifiable
Fee coverage collapses; > 80% speculative	Phase II $\rightarrow$ I regression

Table 27.3: Failure gates and consequences.

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

Proof system failures. A SNARK/STARK is broken or exploited. Mitigation: multi-ZK support; tagging proofs with system/parameter IDs; migration paths.

Hardware capture. A major fab or TEE platform has a backdoor. Mitigation: Layer-0 profiles, lot sampling, open hardware alternatives.

Protocol bugs. Consensus bugs, bridge flaws, privacy leaks. Mitigation: reference implementations, staged rollouts, incident response.

28.2 Economic Risks

Verification cost creep. If verifying proofs ceases to be much cheaper than producing them, the entire asymmetry collapses. Shows up as rising $r(W)$ in VERIFYPRICE.

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.

28.3 Political Risks

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 “na-

tional 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.

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 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**:

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.

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.

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.”

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.

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.

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.

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.

Red Line 8: Governance Capture

Condition: Governance can alter issuance schedules, fee routing, or telemetry rules without hard constitutional constraints, timelocks, or supermajority requirements.

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.

Red Line 9: Wrapper Dominance Becomes Monetary Substitution

Condition: Wrapper Dominance Ratio (Section 11.12) rises for ≥ 2 consecutive quarters while native fee share, private settlement volume, and collateral lockups stagnate or decline. In its full form: custodial and synthetic exposure (CCR, SER) becomes the dominant form of ownership for a sustained period while protocol-native fees, burns, collateral, private settlement, proof demand, and verified-compute usage remain stagnant or declining, with a persistently positive Wrapper–Native Growth Gap (Section 24.8).

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 a threshold share (e.g., $> 50\%$) of VERIFYPRICE-tracked verified compute capacity, i.e., the Enclosure Risk Flag (Section 5.6) triggers and persists.

Why it kills the thesis: AI Money becomes a cloud IOU; the frontier is enclosed rather than homesteaded.

Red Line 12: Agency Failure

Condition: Forced Disclosure Incidence (Section 25.1.2) becomes systemic, or participation/agency use cases (Section 3.4) 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** a threshold share (e.g., $> 50\%$) 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 → proving cost and siting → Physical VERIFYPRICE → 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.

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.

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.

28.6 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.

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.

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 empirical material here draws on energy-sector analysis, principally Doomberg's *Fire Horse* presentation, and on macro work by Luke Gromen. Neither is peer-reviewed; the first is an analyst presentation and the second is commentary. Figures cited are *their* accounting, not independent measurement, and are used to characterize a strategy rather than to establish a quantity.

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 28.5](#) 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.

This is the same structure as the option value argued in [Section 15.9](#): 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 Inversion: Who Is Sovereign

The convergence is structural. The politics are opposite, and the difference is not rhetorical—it appears in specific mechanisms.

Dimension	Closed sovereign stack	Open stack (this thesis)	Mechanism here
Primary sovereign	The state	The participant	Item 9
External trust	Minimized by state capacity	Minimized by verification	Chapter 15
Internal trust	Concentrated at the center	Replaced by public checking	VERIFYPRICE
Privacy	Visible to authorized institutions	Private by default, selective disclosure	Viewing keys, Layer 5
Identity	Administrative and persistent	Contextual and proof-based	Layer 3, PIDL
Settlement	Permissioned, policy-responsive	Non-custodial, bearer-like	Layer 5, PRK
Compute	Nationally integrated	Portable, independently verifiable	Layer 4
Redundancy	Controlled from the center	Produced across independent operators	$\mathcal{O}_s$, Section 15.8
Chief danger	Administrative enclosure	Fragmentation and weak value capture	Lemma 11.1

Table 30.1: Closed and open sovereign stacks: convergent structure, inverted locus of sovereignty.

Read the last row carefully. 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. Both columns have a way of losing, and the thesis is not entitled to compare its best case against the other column's worst.

30.4 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 15.9](#) 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.

30.5 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 3.4](#) 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.

30.6 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.

30.7 What This Changes in the Stack

A chapter that only reframed would not belong in Part VI. Three concrete consequences follow.

1. **The competitor is not fiat.** 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. That competitor wins on convenience and capability, not on principle, and the thesis must therefore compete on exit rights and value capture rather than on capability alone.
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 3.4](#), Red Line 12), not only performance. A scoreboard that measures only capability cannot distinguish the two columns of [Table 30.1](#)—which is the most important thing it could tell us.

Chapter 31

Objections & Responses

Any thesis that ends with “this becomes money” deserves hard pushback.

The objections below are grouped into four 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. 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 four below succeeds, the stack may still be worth building, but it is infrastructure rather than money.

31.1.1 “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.

31.1.2 “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 thesis’s answer is the [Value Capture Lemma](#): the native asset becomes money-like *only if* it is the required fee medium, fees produce burns and staking yield, operators must post it as collateral, issuance is capped, and users cannot bypass it.
- If those conditions fail, the thesis explicitly concedes: the system is useful infrastructure, not a store of value. This is the “utility-token trap” and the thesis names it, not as a risk to dismiss, but as the central condition to satisfy.
- The telemetry regime ([Chapter 24](#)) makes this testable. If native-asset fee share declines while triad usage grows, the objection wins.

31.1.3 “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 28.5](#)).
- 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.

31.1.4 “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 verifiabil-

ity 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 capacity, or energy. The point of PoUW is not to justify waste; it is to redirect energy into **useful work** whose receipts the world must keep buying.
- 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 doing strictly better than random hashing per unit of trust delivered.

- **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 15.9.1). 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 30.7). 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 “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* assets with these properties earn a premium—measurable through fees, burns, collateral, settlement health, and verification cost. Section 28.5 lists thirteen 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 28.5: 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.

31.3.6 “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 persistently above threshold, Red Line 11 (Section 28.5) triggers and the thesis concedes the enclosure path is winning.

31.3.7 “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 11.12). 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 28.5) 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 institutional 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 11.13). 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.

Chapter 32

Why These 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 value-capture conditions hold (non-bypassable fee medium, burns, collateral, issuance discipline), then capacities that deliver Privacy, Proofs, and Compute cheaply and verifiably **may earn monetary premium**—not by decree, but by satisfying the conditions that make an asset worth holding through cycles of repression and technological change.*

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, verified capacity can support markets with fee flows.
4. Fee flows can support a scarce asset *only under non-bypassable accrual rules*.
5. If that asset remains liquid, neutral, verifiable, and legally holdable, it may earn a store-of-value premium.
6. If any condition fails, the system may be useful infrastructure but not money.

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.

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 meet that brief, but in a way very different from gold or Bitcoin:

- **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 will likely choose **stealth default**—negative real yields aided by regulation—over explicit default, a durable SoV must be duration-neutral, peg-proof, and paid for by recurring, indispensable utility. Privacy, proofs, and compute meet that brief: issuance that can’t be decreed; revenues that reprice with fiat budgets; verification that stays cheap and public. **Bonds can be anesthetized; utility cannot.**

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 wrappers around claims on this supply chain. When you hold Work Credits, you hold future access to this stack.

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**. Every monetary epoch can be read as an answer to the question: who has agency over their own savings, payments, and economic identity? Under gold, agency required physical possession. Under fiat, agency was delegated to institutions. Under Bitcoin, agency returned to individuals through self-custody and cryptographic keys. The triad extends this: agency over not just money but over verification (proofs), privacy (settlement without surveillance), and productive capacity (compute without permission). The store-of-value premium accrues to assets that preserve this agency when every other system is designed to extract it.

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—such that Privacy, Proofs, and Compute behave like a next-generation store of value.

That mandate breaks into three concrete programs:

1. **Monetary program.** Design and issue Work Credits so that claims on triad capacity are credibly scarce, easy to verify, hard to seize, and worth holding through repression cycles.
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

A decade ago, Bitcoin taught us that **digital scarcity alone** can be money. The next epoch adds **native utility that the world must keep buying**, secured by proofs anyone can check.

Privacy will preserve agency.

Proofs will anchor truth in a synthetic world.

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 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 condensed geology. Bitcoin condensed thermodynamics. The triad condenses **verifiable agency**: the ability to pay, prove, and compute under adversarial conditions.

This closes the circle opened in [Section 3.4](#). 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 newly monetary. Gold, silver, Bitcoin, and hard assets remain the bridge ([Sections 27.1](#) and [27.1.1](#)). Privacy, Proofs, and Compute become candidate monetary primitives only when the stack proves—through VERIFYPRICE, VERIFYREACH, VERIFYSETTLE, Facility Capacity Receipts, and value-capture telemetry—that it remains private, verifiable, reachable, non-custodial, non-bypassable, and **useful under adversarial conditions without turning its users into dossiers.**

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.

33.4 Two Loops, Not One

One last 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 properties 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.

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.

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.*

Gabaix, X., & Koijen, R. S. J. (2021). In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. NBER Working Paper No. 28967.

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.

Gromen, L. *The Forest for the Trees*. The framing of sovereign debt as a claim on future real surplus, and of monetary credibility as jointly determined by productive capacity and the state’s ability to mobilize it, informs [Section 30.1](#). Commentary rather than peer-reviewed research; used for framing, not for quantities.

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. The sovereign optionality index of [Section 15.8](#) adapts the concept from state-level energy analysis to facility and network scale, dropping the industrial-throughput term as inapplicable and computing the remainder from Facility Capacity Receipt fields. Errors in that adaptation are ours.

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.

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.

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 $\text{PROOF}_{2^{20}}$, 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:

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

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:

$$\text{ProduceProfile}(W) \equiv (p_{50,t}^{\text{prod}}(W), p_{95,t}^{\text{prod}}(W), p_{50,c}^{\text{prod}}(W), p_{95,c}^{\text{prod}}(W))$$

Verification overhead $r(W)$ We define separate time- and cost-based overhead ratios:

$$(v/p)_W^{\text{time}} = \frac{p_{50,t}(W)}{p_{50,t}^{\text{prod}}(W)}, \quad (v/p)_W^{\text{cost}} = \frac{p_{50,c}(W)}{p_{50,c}^{\text{prod}}(W)}$$

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 win-

dow), 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 opera-

tional 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 monetary-grade only if both $\text{VerifyPrice}(W)$ and $\text{PhysicalVerifyPrice}(F)$ clear their respective SLOs; see Section 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 → 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 → 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 — the three cryptographic capacities that can function as monetary primitives.

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 central bridge argument: demand for triad capacity creates SoV premium for the native asset only if five conditions hold—required fee medium, supply reduction via burns, collateral lockup, issuance discipline, and non-bypassability. See [Lemma 11.1](#).

$r(W)$ Verification overhead ratio — $v(W)/p(W)$, where v is verification cost and p is production cost.

Work Credit (WC-Base) The primary monetary object: a scarce native asset required for fees, staking, collateral, and governance. Minted against verified work under issuance constraints. Distinguished from WC-Vouchers (prepaid capacity hedges, not SoV candidates) and Work Receipts (evidence of past work, not scarce).

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.

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](#).

Physical VerifyPrice Used in two related senses, distinguished in [Section 20.8](#): (i) the consti-

tutional 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](#).

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 15.9.2](#).

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 Instruments primarily referencing Privacy + Proofs (shielded settlement capacity, ZK proof capacity).

Attestation Money Instruments backed by proof capacity and provenance attestations. Formerly “Truth Money” in earlier drafts; renamed to emphasize that proofs verify bounded claims (origin, computation, custody, policy compliance), not semantic truth.

AI Money Instruments primarily referencing Compute + Proofs (verified FLOPs, inference capacity).

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 11.13](#).

Native instrument A protocol-internal claim (Work Credits, staking or LP positions, corridor claims). Using it exercises fees, burns, or collateral.

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. $\varepsilon \approx 0$ means holders sit still; $\varepsilon \approx -1$ means they trim to a constant dollar position. Measured, not structural. See [Appendix H.1](#).

Net mechanical gain (κ) $L(L-1) + \varepsilon L^2$; the coefficient converting an underlying return into a wrapper’s total return-coupled exposure demand. See [Appendix H.3](#).

Wrapper Recycling Ratio (WRR) $-\varepsilon L/(L-1)$. 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.6](#).

Wrapper–Native Growth Gap (WNG) $\Delta \ln E^{\text{wrapper}} - \Delta \ln U^{\text{native}}$. Persistently positive means financial exposure is outgrowing monetary use.

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](#).

Appendix H

Formal Model of Market Realization, Wrapper Flows, and Price Capture

This appendix develops the flow machinery behind the Market Realization Plane (Section 11.13) 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 follows the inelastic-markets literature, especially Gabaix and Koijen (2021). 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, $\text{NAV}_{i,t}$ 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 must be published rather than absorbed into either. Assigning α_i to Q^{RC} would overstate momentum amplification; discarding it would understate the structural bid.

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.

H.6 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 :

$$Q_{j,t}^{RD} = \sum_k w_{jk,t} F_{k,t}$$

where $F_{k,t}$ is net creation flow 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 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.7 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.8 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 ra-

tio. 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 7.3](#): 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.9 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.10 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.
- 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.11 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.12 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 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.