



Get Paid  USDT Every Time, In Real-Time

CONTENT REWARDS

Whitepaper: An automated system that manages creator relationships and payments

• Last Updated March 20 2026



Contents

1	Abstract	5
2	The Limitations of Content Rewards on Whop Now	9
3	PropWhop: A Proprietary Solution for Content Rewards on Whop	10
3.1	Overview	10
3.2	Aptos	10
3.3	Shelby	10
3.4	The Integration	11
4	Where the Only Thing That Matters Is Your Results	11
4.1	Earn Without Permission	11
4.2	Proof-of-Results as the Onramp	11
4.3	Offramp Friction as a Whop Retention Flywheel	11
5	Content Origin Verification	12
5.1	Why Origin Verification Exists	12
5.2	How Shelby Stores and Fingerprints Source Material	12
5.3	Audio and Visual Sliders	13
5.4	The Confidence Score Formula	14
5.4.1	Worked Examples	15
5.5	Multi-Source Clips	16
5.6	Eligibility: The 80% Confidence Floor	17
6	Campaign CPM Curve and the Content Rewards Trust Score	18
6.1	How the CPM Range Works	18
6.2	How a Pricing Model Eliminates Gatekeeping	18
6.3	How the Trust Score Prevents Botting	18
7	The Two-Dimensional CPM Model	18
7.1	Anti-Bot CPM Floor and Campaign Hijack Prevention	19
8	Submission Similarity and Account Scoring	21
9	From Manual Gatekeeper to Automated Campaign Operator	21
9.1	Why Manual Arbitration Is the Bottleneck	21
9.2	Why Shelby and Aptos Fix This	22
10	Real-Time Clipper Eligibility and Payouts	22
10.1	The Problem With the Old System	22
10.2	How Programmatic Verification Solves This	23
11	Account Farming and Acceptable Fraud Tolerance	23
11.1	The Attack	23
11.2	Account Reselling Detection	23
11.3	Graduated Scoring Over Bans	23
11.4	The Acceptable Outcome Is ROI Focused	24
12	Phase 1: Objective-Only Requirements	24
12.1	What Phase 1 Means	24
12.2	Why Objective-Only at Launch	24

- 13 Campaign Creation Fee Flow** **25**
- 13.1 The 20% Surcharge 25
- 13.2 Why the Creator Pays, Not the Clipper 25
- 13.3 Whop: 10% Fee 26
- 13.4 Content Rewards: 7% Fee 26
- 13.4.1 Dynamic Utilization Curve 26
- 13.4.2 CRIF Overflow to WLP 27
- 13.5 Aptos: 3% Fee 27
- 14 Perpetual Campaign** **28**
- 14.1 Dynamic APT Budgeting 28
- 15 Creator Trust Score** **29**
- 15.1 What the Creator Trust Score Encodes 29
- 15.2 The Three Dimensions 29
- 16 Clipper Trust Score** **29**
- 16.1 What the Clipper Trust Score Encodes 29
- 16.2 What Feeds the Score 30
- 16.3 How the Score Affects Earnings 30
- 17 Campaign Trust Score** **30**
- 17.1 What the Campaign Trust Score Encodes 30
- 17.2 Campaign Trust Score Inputs 30
- 18 Campaign Lifecycle Actions** **32**
- 18.1 Campaign Optimization (Most Common) 32
- 18.2 Sunset Campaign (Graceful Wind-Down) 32
- 18.2.1 Sunset Haircut Formula 33
- 18.3 End Campaign (Immediate) 34
- 18.4 Dispute Campaign (Last Resort) 35
- 19 Bounty Hunters** **36**
- 19.1 Compensation Model 36
- 20 Automated Dispute Resolution (ADR)** **36**
- 20.1 The ADR Flow and Timeline 37
- 21 Content Rewards Insurance Fund (CRIF)** **38**
- 21.1 Insurance Reserve 38
- 21.2 Lien Escrow 39
- 22 Whop Liquidity Provider (WLP)** **39**
- 22.1 Clipper Earnings and Automatic WLP Deposit 39
- 22.1.1 Gasless UX and Cost Sustainability 40
- 22.2 Withdrawability 40
- 22.3 Withdrawals 40
- 22.4 Yield, Voting, and Risk Structure 40
- 23 Auto-Deleveraging (ADL) and the Loss Waterfall** **41**
- 23.1 The Backstop Waterfall (Four Tiers) 41
- 23.1.1 Tier 1: Lien Liquidation (Principal and Yield Exposed) 41
- 23.1.2 Tier 2: Insurance Fund Absorption (Insurance Before Socializing Losses) 41
- 23.1.3 Tier 3: Whop Liquidity Provider ADL (Yield-Only, Principal Protected) 42

- 23.1.4 Tier 4: Future Campaign Fee Credits (Final Backstop) 42
- 24 Future Campaign Fee Credits (FCFC) 44**
 - 24.1 How It Works 44
 - 24.2 Decay Mechanics 44
- 25 Creator Dispute Payout Priority (CDPP) 44**
 - 25.1 The Payout Method (Three Tiers) 45
 - 25.1.1 Tier 1: Filing Fees First (Entry Costs Refunded) 45
 - 25.1.2 Tier 2: Pro Rata on Confirmed Claims (All Creators Reimbursed Equally) 45
 - 25.1.3 Tier 3: Future Campaign Fee Credits (Relaunch Campaigns Without Fee) 45
- 26 Roadmap to Phase 2 45**
- 27 Summary 47**

1 Abstract

Whop is in a position remarkably similar to early Shopify, as a platform that has quietly become the infrastructure layer for a new class of internet business, growing at a pace that suggests the market is far larger than current adoption reflects. Shopify gave anyone the ability to open a storefront. Whop gives anyone the ability to sell digital products, communities, and services. The trajectories rhyme: explosive creator onboarding, rapidly expanding GMV, and a platform architecture that is becoming the default operating system for an entire generation of digital entrepreneurs to earn their first dollar online.

What made Shopify's growth era explosive was not just the storefront. It was the combination of Shopify, the Oberlo app, and Facebook Ads, as a system where a creator could source a product, build a store, and acquire customers through paid advertising with trackable ROI, all within an interconnected workflow.

Oberlo solved quality control and supply chain management, connecting reputable manufacturers with store owners without going off platform, within a fully integrated storefront automating order fulfillment and providing tracking, reducing the barrier for new businesses to frictionlessly source products in a growing marketplace, and track everything end-to-end.

Facebook Ads solved distribution and attribution with a self-serve platform that gave store owners granular audience targeting and pixel-based conversion tracking, showing exactly which ad, audience, and creative produced each sale in real time. Merchants could kill losing campaigns within hours and scale winners with confidence. Attribution was what made the entire ecosystem scalable, because without it, merchants were inefficient with budgets and couldn't properly scale.

Shopify solved store creation without development costs, lowering the barrier from hiring a developer to clicking a button. As the merchant base grew, Shopify's earnings diversified beyond transaction fees into app store revenue, tiered subscription plans, and Shopify Capital, transforming the platform from a storefront provider into a resilient commerce operating system.

The flywheel was self-reinforcing: more stores meant more ad spend, more ad spend meant more data for Facebook's targeting algorithm, better targeting meant higher ROI for merchants, higher ROI meant more stores, and more stores meant more apps, more subscriptions, and more capital deployed through Shopify's expanding ecosystem. That combination was the face of the internet from 2017 onwards, where Shopify went from a few billion dollar valuation to hundreds of billions, with many millionaires created as a byproduct.




Figure 1: Shopify × Whop — The Oberlo Parallel

Content Rewards maps directly onto this model for Whop:

Where Oberlo connected merchants with manufacturers and automated fulfillment end-to-end, Content Rewards connects campaign creators with clippers and automates the entire workflow from source material delivery to verified payout. Creators supply their existing content as campaign assets, and clippers produce and distribute short-form derivatives across TikTok, Instagram, and YouTube without anyone leaving the platform.

Where Facebook Ads gave merchants self-serve distribution with pixel-based attribution to measure exactly what was working, Shelby's read-metering and verification stack give campaign creators independent, granular proof of which clips generated real views from real audiences. This is the attribution layer that lets them scale winning campaigns and cut losing ones with the same confidence.

Where Shopify evolved from a storefront into a full commerce operating system, Whop is positioned to follow the same trajectory, with Content Rewards as the initial growth engine and ecosystem expansion into finance, app store revenue, and creator services compounding behind it.

Paid advertising delivers more consistent results by allowing creators to scale through budget increases, but simply adding a zero or two to the budget doesn't mean ROI performance will scale proportionately. Ad fatigue exhausts audiences, promotional content is approached with increasing skepticism, and the cost per acquisition rises as the easy conversions are captured first. The creator pays before value is delivered, and returns diminish as spend grows.

Short-form content is maximally efficient per dollar of spend because it's organic. Cheaply produced clips can drive massive results in a single spike, but that spike is the downfall: results are inconsistent, hard to predict, and nearly impossible to scale with budget increases. A creator can't reliably turn \$1,000 of clipping spend into \$10,000 of clipping spend and expect proportional returns, because virality is not a linear function of investment.

Content Rewards on Whop, with the help of Shelby and Aptos, takes the benefits of both without the drawbacks of either. The scalability and budget control of paid advertising (creators set CPM ranges, fund campaigns with defined budgets, and track ROI through verified attribution) combined with the organic efficiency and audience trust of short-form content (clips live natively on social platforms as content people actually want to watch,

not ads they instinctively scroll past). The creator pays after value is delivered, not before. Results scale with budget because more budget attracts more clippers producing more clips across more platforms, and the per-view economics remain stable because the CPM model prices each view independently rather than degrading with volume.

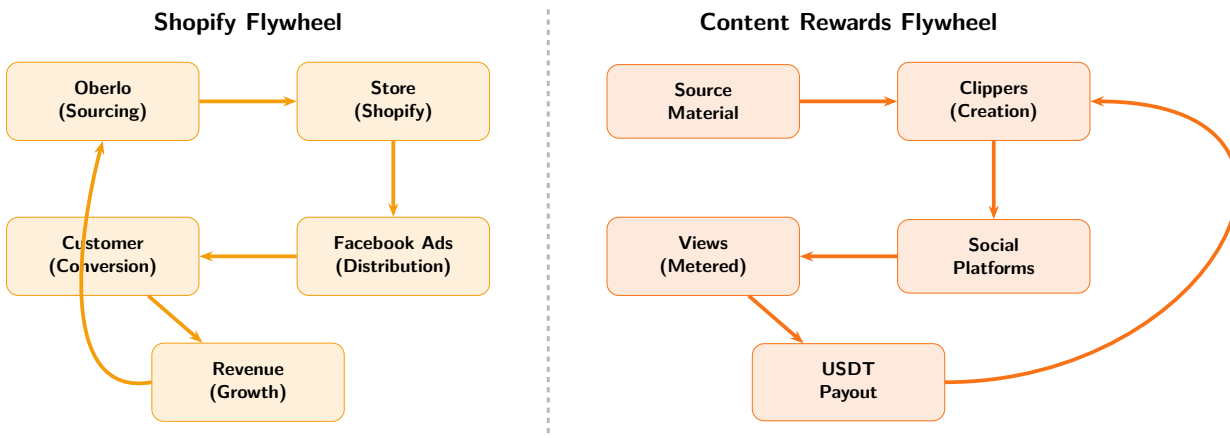


Figure 2: Dual Flywheel: Shopify’s proven model mapped onto Content Rewards.

The Content Rewards flywheel mirrors Shopify’s proven growth engine: automated sourcing (Oberlo → clippers), measurable distribution (Facebook Ads → Shelby read-metering), and instant revenue recognition. PropWhop is the infrastructure that makes this flywheel spin without friction.

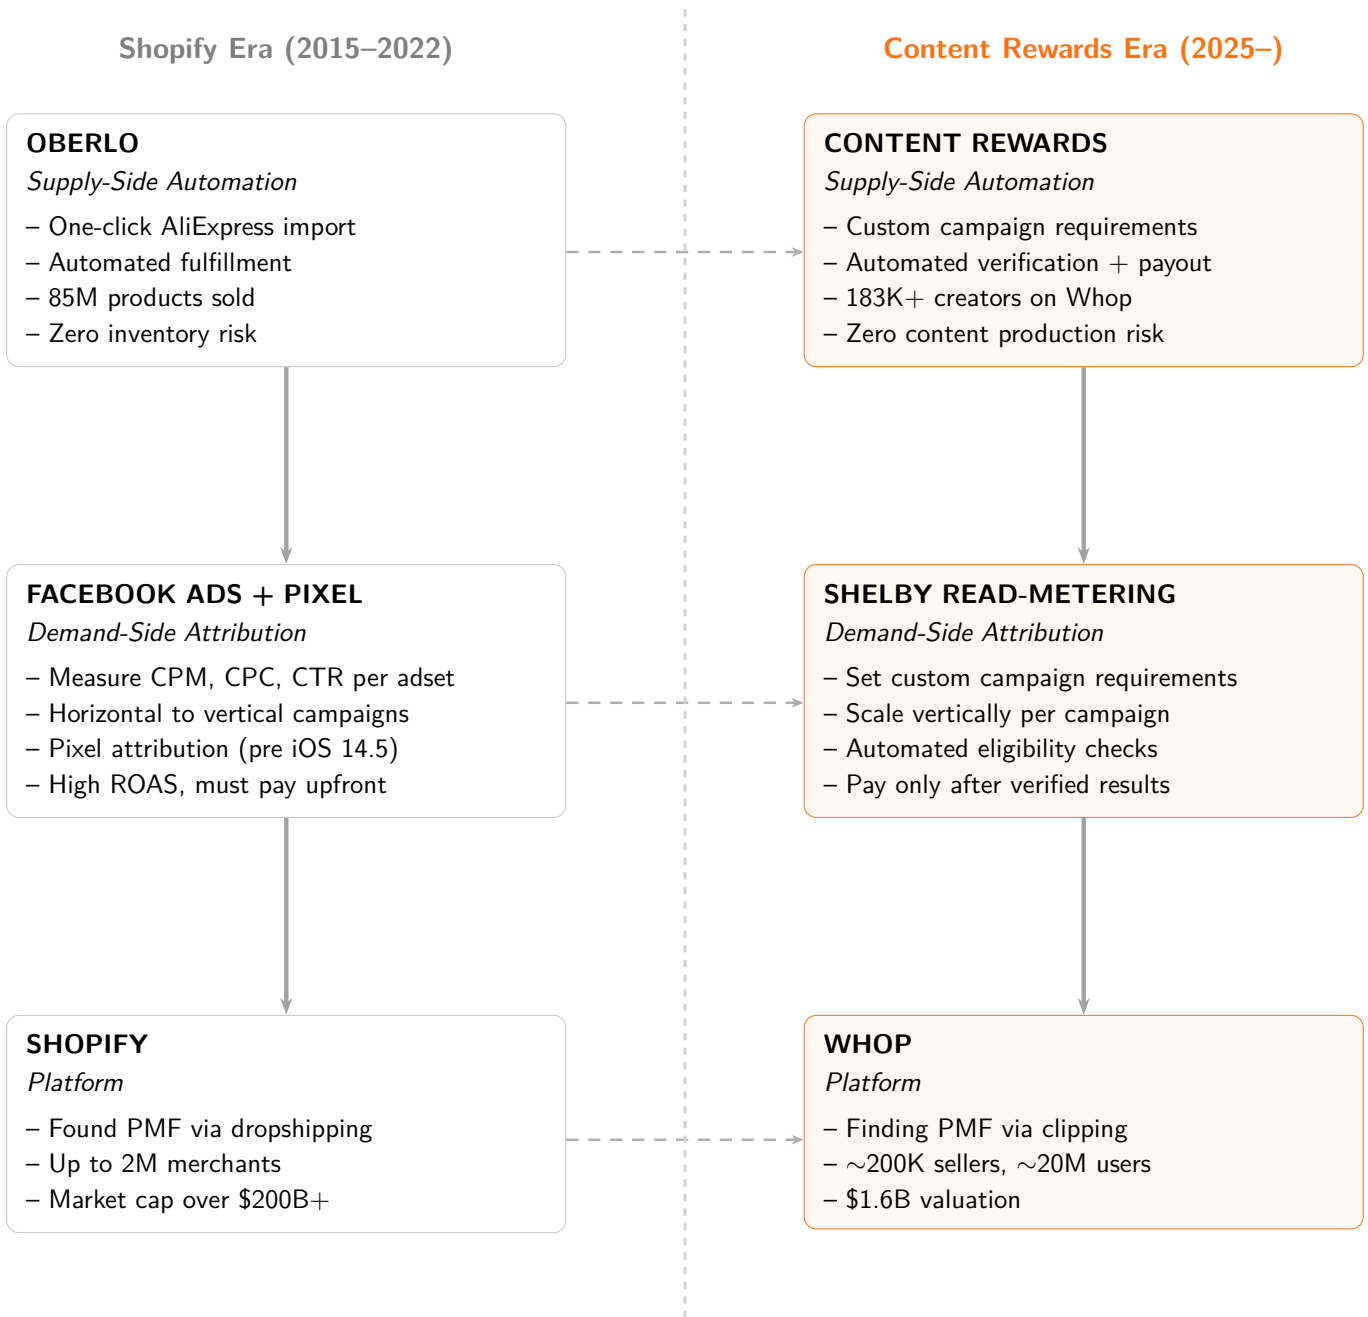


Figure 3: The structural parallel: Content Rewards combines Oberlo (supply automation) and Facebook Ads (demand attribution) into a single product for Whop.

The Key Difference

Oberlo’s flywheel was **killed by iOS 14.5** — Apple’s privacy change destroyed Facebook Pixel attribution, collapsing ROAS by 38% and costing Meta \$10B in 2022 revenue. Content Rewards’ attribution is **cryptographic and platform-independent**. Shelby read-metering does not depend on cookies, device IDs, or tracking pixels. It cannot be disrupted by a privacy policy change.

2 The Limitations of Content Rewards on Whop Now

The problem with the current Content Rewards model on Whop, and the reason it has not yet reached escape velocity, is that every layer of the system — verification, payouts, eligibility, fraud prevention, and dispute resolution — relies on manual processes that break under volume and create misaligned incentives between creators and clippers.

- View counts are self-reported with no independent verification, and eligibility decisions are subjective where the campaign creator is both the rule-setter and the judge, able to set vague requirements like "must be high quality" and selectively reject compliant submissions.
- Payouts are delayed behind manual review queues that become a lottery when budgets run low or clips go viral — two clippers submit equally eligible content on the same day, one gets reviewed and paid, the other sits in queue until the budget is gone, denying payment to the clippers creating the most value at exactly the moment the system is working as designed.
- Fraud detection relies on human review that scales linearly with volume. Disputes escalate from the creator's inbox, to Content Rewards, to Whop's inbox, where humans with imperfect information make judgment calls or send users in loops without proper adjudication — more campaigns means more disputes means more support staff means higher costs, and at scale the system breaks from operational stress.
- A creator receiving many submissions spends hours on manual review, inflating costs with virtual assistants who make inconsistent decisions. The creator's time and money goes to back-office disputes and payroll admin, instead of creating more content that clippers can use to go viral.
- Neither side can trust the system enough to commit. Creators cannot scale budgets aggressively with long holds, subjective rejections, and unverifiable metrics. Clippers cannot treat it as reliable income when payout unpredictability is the core reason talent churns. Trust erodes on both sides, and the flywheel dies at the trust layer.

The Structural Problems

1. **View counts are unverifiable** — self-reported metrics with no independent verification; eligibility decisions are subjective where the campaign creator is both rule-setter and judge.
2. **Payments are delayed and manual** — review queues become a lottery when budgets run low; two equally compliant submissions on the same day, one gets paid, the other does not.
3. **Fraud detection does not scale** — human review grows linearly with volume; more campaigns means more disputes means more support staff means higher costs.
4. **Creator time is wasted on review** — hours spent manually reviewing submissions, inflating costs with virtual assistants who make inconsistent decisions.
5. **Budget exhaustion is unpredictable** — clippers generating the most value get denied payment by queue timing at exactly the moment the system is working as designed.
6. **Dispute resolution is a loop** — escalates from creator inbox to Content Rewards to Whop, where humans with imperfect information make judgment calls or send users in circles.
7. **Creator incentives misalign at scale** — neither creators nor clippers can trust the system enough to commit; trust erodes on both sides.
8. **The whole system relies on trust, not math** — no cryptographic guarantees, no on-chain verification; every layer breaks under volume and creates misaligned incentives.

Let's fix it.



3 PropWhop: A Proprietary Solution for Content Rewards on Whop

This documentation describes a proprietary solution that replaces every manual component with programmatic verification, real-time settlement, and an economic structure that makes fraud unprofitable by design. The system is built on infrastructure with direct lineage to Facebook’s own attempts to solve payments and content distribution at global scale, technology that was engineered to serve billions of users and is now applied to the specific problem of making Content Rewards trustworthy enough that campaign creators scale their budgets with the same confidence that Shopify merchants scaled their Facebook Ad spend. Aptos and Shelby fix this by providing trustless infrastructure to drastically improve content eligibility criteria in a verifiable way accelerating payouts to clippers, while also automating dispute resolutions quickly with market-driven incentives instead of creating backlogs and support tickets.

3.1 Overview

With the help of Shelby and Aptos, Content Rewards truly can be an automated marketing system on Whop that connects brands and creators with content clippers, where brands fund campaigns, clippers create and distribute short-form content derived from the brand’s source material, and clippers are instantly paid for their virality results based on verifiable criteria, earning per eligible view in real time.

3.2 Aptos

The system is built on Aptos, a Layer 1 blockchain created by the core team that designed Facebook’s Libra and Diem payment infrastructure. When regulators killed Diem, the team launched Aptos independently with the open-source codebase, backed by over \$400 million from institutional investors. The total investment into the infrastructure that underpins PropWhop exceeds \$1.4 billion across over eight years of research and development — a technology stack purpose-built at Facebook scale for money movement at its core, that Whop and Content Rewards can plug into directly without bearing additional developer workload.



Figure 4: Shelby and Aptos: decentralized hot storage with native read-metering (Shelby) and sub-second on-chain settlement (Aptos).

3.3 Shelby

The decentralized hot storage protocol built by the same team, engineers who scaled content delivery across Meta’s consumer platforms including Instagram — infrastructure that served billions of reads per day. Shelby applies that experience to a problem Instagram never solved: making every content view independently verifiable and monetizable. The protocol’s native micropayment metering means every read automatically generates timestamped, geo-tagged, wallet-identified access data — the raw signal that powers Content Rewards’ entire

verification, fingerprinting, and payout stack, not as an analytics layer bolted on after the fact, but as a byproduct of how the protocol moves content. For Content Rewards and Whop, this means every clip view, every source access, and every payout trigger is provable by default — giving both creators and clippers neutral evidence that neither side controls, eliminating the trust gap that manual review could never close.

3.4 The Integration

The integration of these technologies transforms Content Rewards from a manually reviewed, delayed-payout system into a programmatically verified, instantly settled, market-driven infrastructure that scales without proportionally scaling headcount.

Content Rewards, as Whop's product with the strongest existing product-market fit, serves as the natural trojan horse for introducing blockchain-native payments and financial infrastructure to Whop's entire creator base, a strategic entry point following Tether's \$200 million investment. Creators who experience instant, verifiable USDT payouts through Content Rewards become native users of on-chain financial rails without needing to understand the underlying technology.

The detailed mechanisms behind how the verification signals are generated, how they feed into the account scoring system, and how they create the economic structure that makes fraud unprofitable at the margins are described throughout the following sections.

4 Where the Only Thing That Matters Is Your Results

4.1 Earn Without Permission

Content Rewards is how the next generation can earn their first dollar online on Whop, getting paid directly peer-to-peer without KYC. Many clippers are underage and don't have bank accounts, can't open payment processors to receive their payout without using their parents' documentation, adding friction.

The precedent already exists. Before Ethereum's merge to proof-of-stake, Chinese citizens were banned from crypto exchanges and blocked from bank transfers, but mining created an alternative: converting electricity into ETH, bypassing every financial gatekeeper. China dominated global ETH mining hash rate because proof-of-work provided a permissionless path to earn when every other path was closed. When people can convert energy into real money without permission from a bank or government, adoption scales to wherever the opportunity exists.

4.2 Proof-of-Results as the Onramp

Content Rewards creates the same dynamic through proof-of-results. Clippers convert creative energy into USDT with value-add viral clipping. The work is real, the output is verifiable through Shelby, and the payout is direct: USDT streamed without a bank, a payment processor, or an identity check. A teenager who video edits in Berlin, Seoul, or Sao Paulo can start clipping and receive USDT for all eligible views within the same day. The only thing that matters is results. No bank account required unless offramping.

4.3 Offramp Friction as a Whop Retention Flywheel

The same no-KYC property that makes earning frictionless creates natural USDT retention. Offramping to a bank account requires KYC at the bank, crypto debit card, or centralized exchange; identity verification, bank linking, and processing delays. For many clippers, the offramp is the hardest step. The path of least resistance is keeping USDT inside Whop, where earnings can generate yield in the WLP, circulate across Whop Finance ecosystem, be spent on other Whop products and services, or be reinvested into their own community. Every transaction within Whop is P2P, wallet-to-wallet with no identity gate. More opportunities on Whop, the less reason to offramp, and the more capital circulates internally, earning consistent fees.

5 Content Origin Verification

5.1 Why Origin Verification Exists

Every clip submitted to a Content Rewards campaign must prove it actually uses the brand's source material. Without this, anyone could submit unrelated content to a campaign and claim payouts on views that have nothing to do with the brand. Origin verification is the gate that separates legitimate clipping from campaign hijacking, where the campaign budget is siphoned from the creator with irrelevant content.

5.2 How Shelby Stores and Fingerprints Source Material

When a brand uploads source material (podcast episodes, YouTube videos, lifestyle vlogs, fitness content, travel footage) to a Content Rewards campaign, Shelby stores the original files as Blobs with cryptographic commitments on Aptos. Alongside the source material, the system generates and stores two fingerprint databases:

Audio fingerprints: The source audio track can be divided into 1–1.5 second windows with 75% overlap, producing a new fingerprint roughly every 0.25–0.375 seconds. Each window receives a perceptual hash, a compact representation of the spectral content that survives compression, level changes, speed adjustments, and background music additions. The sub-second resolution is necessary because clippers cutting between multiple source assets transition in under a second; coarser windows would span two sources and degrade match accuracy at the seam. A 2-hour podcast produces roughly 19,000–28,000 fingerprint windows stored as a separate Shelby Blob. Each hash is 32–256 bytes, so the total storage overhead is modest with Shelby erasure coding.

Visual fingerprints: Keyframes can be extracted from the source video at 2–4 frames per second. Each keyframe receives a perceptual hash that captures structural visual content, robust to resolution changes, aspect ratio cropping, caption overlays, color grading, and zoom effects. The higher frame density ensures that transitions, hook intros, and creative inserts each get their own frames rather than being absorbed into adjacent source frames. A 30-second clip at 3 fps produces 90 keyframes, enough that a 0.5-second transition only affects 1–2 frames instead of dominating the match calculation. These are stored as another Shelby Blob alongside the source.

When a clipper submits a clip, the system extracts audio and visual fingerprints from the submission and matches them against the source databases. The output is a Content Origin Report containing:

- The percentage of submission audio that matches source audio (audio match)
- The percentage of submission visual keyframes that match source frames (visual match)
- A timestamp mapping showing exactly which segments of the source appear in the clip
- Identification of non-source additions (hooks, transitions, B-roll, outros)

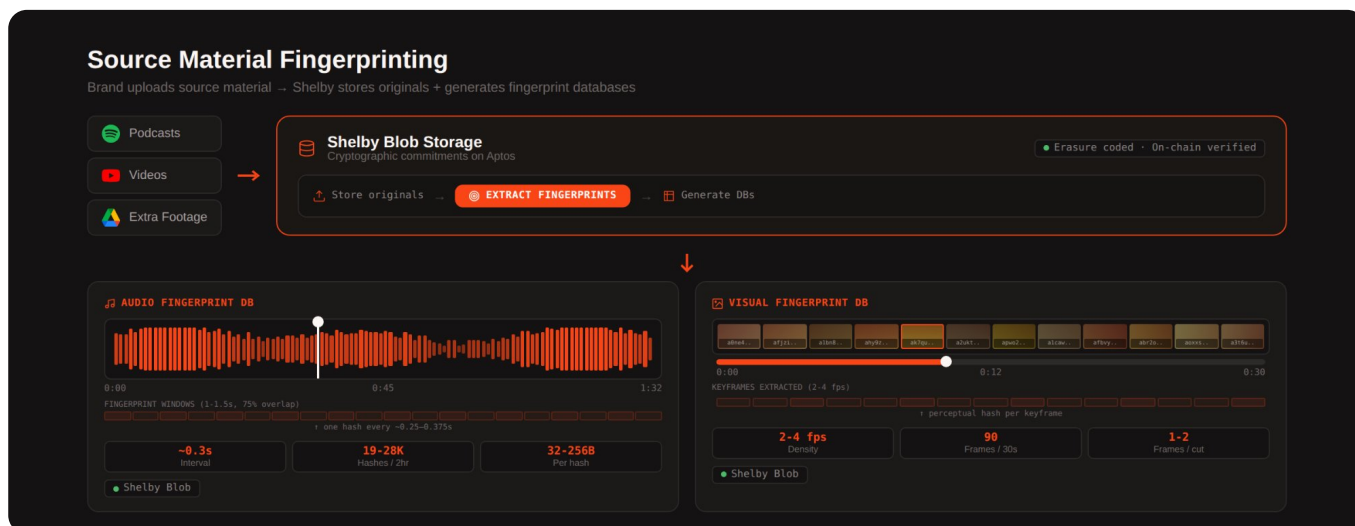


Figure 5: Source Material Fingerprinting: dual-channel audio and visual perceptual hashing stored as cryptographically committed Blobs on Shelby.

5.3 Audio and Visual Sliders

Different source content delivers value through different channels. A podcast is audio-first, where the creator’s voice is the content and B-roll visuals enhance the clip but aren’t the value driver. Lifestyle content (fitness routines, travel vlogs, day-in-the-life, cooking, fashion) is visual-first, where what the viewer sees is the content and the audio can be enhanced with motivational overlays or trending sounds. A live performance or informational vlog delivers value through both channels equally.

Each source asset in a campaign gets two independent sliders that tell the system how much each channel matters for that asset:

Audio slider (0–100): How much the brand values original source audio in clips. - 0 = Audio fidelity doesn’t matter. Clippers can replace the audio entirely with voiceovers, trending sounds, or music. Appropriate when the brand’s value lives entirely in the visuals. - 100 = The brand wants the exact original audio preserved. Every word of the creator’s voice matters. Appropriate for podcasts, interviews, and spoken-word content where the creator’s voice IS the product.

Visual slider (0–100): How much the brand values original source visuals in clips. - 0 = Visual fidelity doesn’t matter. Clippers can use B-roll, stock footage, or their own visuals. Appropriate when the brand’s value lives entirely in the audio. - 100 = The brand wants the exact original visuals preserved. Appropriate for lifestyle, fitness, and travel content where what the viewer sees IS the product.

The sliders are independent; they don’t need to sum to any number. A brand can set: - Audio 90, Visual 20 (podcast: keep my voice, visuals are flexible) - Audio 15, Visual 85 (lifestyle vlog: keep my footage, audio overlay is fine) - Audio 70, Visual 70 (live performance: both matter equally)

The table below provides recommended starting points by content type. These are guidelines, not requirements — creators adjust to match their specific source material.

Content Type	s_a	s_v	$w_a : w_v$	Rationale
Podcast / Inter-view	80	20	80 : 20	Voice is the product. B-roll, captions, and visual overlays are expected and encouraged.
Lifestyle / Fitness / Travel vlog	20	85	19 : 81	Footage is the product. Trending audio overlays and music swaps are fine.
Live performance	70	70	50 : 50	Both channels carry equal value. Neither can be heavily modified.
Informational / Tutorial	60	60	50 : 50	Balanced — the explanation (audio) and the demonstration (visual) are both essential.
Music / Audio-first	95	10	90 : 10	The track is everything. Visuals are just a vehicle for distribution on short-form platforms.

Table 1: Slider configuration reference by content type. s_a = audio slider, s_v = visual slider, $w_a : w_v$ = effective weight ratio after normalization. Minimum combined slider value: $s_a + s_v \geq 80$.

5.4 The Confidence Score Formula

The confidence score is a weighted average. The sliders set the weights. No curve, no nonlinearity, just a direct measurement of how well a clip matches the brand's priorities.

Step 1: Normalize the sliders into proportional weights.

$$w_{\text{audio}} = \frac{\text{audio_slider}}{\text{audio_slider} + \text{visual_slider}} \quad w_{\text{visual}} = \frac{\text{visual_slider}}{\text{audio_slider} + \text{visual_slider}}$$

This normalization ensures the confidence score always falls on a 0–100 scale, regardless of where the brand sets the sliders. Audio 80 / Visual 20 produces the same relative weights as Audio 40 / Visual 10; in both cases, audio is 4x more important than visual. What matters is the ratio between the sliders, because the ratio defines how much each channel contributes to the overall score.

Step 2: Compute the weighted average of the two channel matches.

$$\text{confidence} = (\text{audio_match} \times w_{\text{audio}}) + (\text{visual_match} \times w_{\text{visual}})$$

The confidence score is a single number between 0 and 100 that represents how well the clip satisfies the brand's channel priorities for that source asset.

● WEIGHT NORMALIZATION

$$w_{\text{audio}} = \frac{S_{\text{audio}}}{S_{\text{audio}} + S_{\text{visual}}}$$

$$w_{\text{visual}} = \frac{S_{\text{visual}}}{S_{\text{audio}} + S_{\text{visual}}}$$

Figure 6: Weight Normalization: audio and visual sliders normalized to weights summing to 1.0.

● CONFIDENCE SCORE

$$C = (M_{\text{audio}} \times w_{\text{audio}}) + (M_{\text{visual}} \times w_{\text{visual}})$$

Figure 7: Confidence Score formula: weighted sum of audio and visual match percentages.

5.4.1 Worked Examples

Podcast | Audio Slider: 80 | Visual Slider: 20

The brand's voice is the product. B-roll is fine. Replacing the audio isn't.

$$w_{\text{audio}} = \frac{80}{80 + 20} = 0.80 \qquad w_{\text{visual}} = \frac{20}{80 + 20} = 0.20$$

Clip Type	Audio	Visual	Score	Result
B-roll heavy clip	95%	30%	$(95 \times 0.80) + (30 \times 0.20) = 82.0\%$	✓ Pass
Raw clip	99%	98%	$(99 \times 0.80) + (98 \times 0.20) = 98.8\%$	✓ Top CPM
Audio overlay	40%	95%	$(40 \times 0.80) + (95 \times 0.20) = 51.0\%$	✗ Fail

Lifestyle Vlog | Audio Slider: 20 | Visual Slider: 85

The creator's footage is the product. Trending audio overlays are fine. Replacing the visuals isn't.

$$w_{\text{audio}} = \frac{20}{20 + 85} = 0.19 \qquad w_{\text{visual}} = \frac{85}{20 + 85} = 0.81$$

Clip Type	Audio	Visual	Score	Result
Trending audio overlay	45%	92%	$(45 \times 0.19) + (92 \times 0.81) = 83.1\%$	✓ Pass
Replaced visuals, original audio	96%	25%	$(96 \times 0.19) + (25 \times 0.81) = 38.5\%$	✗ Fail

Live Performance | Audio Slider: 70 | Visual Slider: 70

Both channels carry the value equally. You can't heavily modify either one.

$$w_{\text{audio}} = \frac{70}{70 + 70} = 0.50 \qquad w_{\text{visual}} = \frac{70}{70 + 70} = 0.50$$

Clip Type	Audio	Visual	Score	Result
Clean clip with captions	93%	88%	$(93 \times 0.50) + (88 \times 0.50) = 90.5\%$	✓ High CPM
B-roll heavy clip	94%	35%	$(94 \times 0.50) + (35 \times 0.50) = 64.5\%$	✗ Fail

5.5 Multi-Source Clips

A clipper may use content from multiple source assets in the same posted clip, for example stitching 60% of the clip from podcast episode 1 and 40% from episode 2. Each episode may have different slider settings.

The sub-second audio windows and 2–4 fps visual sampling described above are what make this possible. The system needs to identify precisely where one source ends and another begins; if the boundary between sources falls inside a coarse fingerprint window, that window partially matches both sources and degrades accuracy for both. Tighter sampling keeps the boundary clean so each asset's confidence is computed only against the frames and audio that actually came from it.

The system handles this in two steps:

Step 1: Compute a separate confidence score for each source asset using that asset's own slider weights and the match percentages for the portions of the clip derived from that asset.

$$c_1 = (\text{audio_match}_1 \times w_{\text{audio},1}) + (\text{visual_match}_1 \times w_{\text{visual},1})$$

$$c_2 = (\text{audio_match}_2 \times w_{\text{audio},2}) + (\text{visual_match}_2 \times w_{\text{visual},2})$$

Step 2: Blend the per-asset confidence scores proportionally by how much of each source appears in the submission.

$$p_1 = \frac{\text{matched_duration}_1}{\text{total_matched_duration}}, \quad p_2 = \frac{\text{matched_duration}_2}{\text{total_matched_duration}}$$

$$\text{overall_confidence} = (c_1 \times p_1) + (c_2 \times p_2)$$

Example: A 45-second clip uses 27 seconds from video 1 (audio 80, visual 20) and 18 seconds from video 2 (audio 30, visual 75):

$$c_1 = 88\%, \quad c_2 = 84\%, \quad p_1 = 0.60, \quad p_2 = 0.40$$

$$\text{overall} = (88\% \times 0.60) + (84\% \times 0.40) = 52.8 + 33.6 = \mathbf{86.4\%} \rightarrow \text{passes 80\% floor } \checkmark$$

Each source asset's slider preferences are respected proportionally. The brand's intent for each piece of content carries through even when a clipper creatively combines multiple sources.

● PER-ASSET CONFIDENCE

$$C_{\text{asset 1}} = (M_{\text{audio}}^{(1)} \times w_{\text{audio}}^{(1)}) + (M_{\text{visual}}^{(1)} \times w_{\text{visual}}^{(1)})$$

$$C_{\text{asset 2}} = (M_{\text{audio}}^{(2)} \times w_{\text{audio}}^{(2)}) + (M_{\text{visual}}^{(2)} \times w_{\text{visual}}^{(2)})$$

Figure 8: Per-Asset Confidence: each source asset scored independently before blending.

● PROPORTIONAL BLENDING

$$p_1 = \frac{d_1}{d_1 + d_2} \quad p_2 = \frac{d_2}{d_1 + d_2}$$

$$C_{\text{overall}} = (C_{\text{asset 1}} \times p_1) + (C_{\text{asset 2}} \times p_2)$$

Figure 9: Proportional Blending: final confidence weighted by duration of each source asset used.

5.6 Eligibility: The 80% Confidence Floor

Eligibility is binary. If the overall confidence score (single-source or multi-source blended) meets or exceeds the campaign's confidence floor, the submission is verified and eligible. USDT streaming begins immediately.

The default minimum floor is 80%. This threshold is calibrated for Phase 1, where the dispute resolution infrastructure (CRIF reserve, WLP liquidity, bounty hunter network) is still building capacity. At 80%, the system optimizes for low dispute volume: the verification is strict enough that the vast majority of eligible submissions are genuinely derived from the brand's source material, and the vast majority of ineligible submissions are correctly rejected.

The brand can raise the floor above 80% (requiring higher fidelity) but cannot lower it below 80%, as a precautionary measure to minimize dispute volume in the beginning of Phase 1.

Content Rewards can enforce a minimum combined slider value: $\text{audio_slider} + \text{visual_slider}$ must be at least 80. This prevents a brand from setting both sliders so low that the confidence formula becomes trivially easy to satisfy.

6 Campaign CPM Curve and the Content Rewards Trust Score

6.1 How the CPM Range Works

The creator sets a floor CPM and a ceiling CPM for their campaign:

- The floor is subject to a protocol-recommended minimum calibrated to make view botting unprofitable for that campaign's specific requirement set. The creator can raise the floor but not lower it below the anti-bot threshold.
- The ceiling is entirely the brand's choice for how much they're willing to pay for the best clippers with the highest-fidelity content and associated track record.

Between the floor and ceiling, the creator selects a CR Curve shape (linear, convex, or concave) that determines how the Trust Score maps to rate, depending on how much they value quality over quantity.

6.2 How a Pricing Model Eliminates Gatekeeping

No applications. Every campaign is open to every clipper. A brand-new account can submit to any campaign on the platform from day one. The mechanism that protects the brand isn't access restriction; it's payout pricing. New accounts earn floor CPM. Proven accounts earn ceiling CPM. Both get a fair deal at the risk-appropriate price.

No private deals. The CPM range gives proven clippers a premium rate that matches or exceeds what they'd negotiate in a private arrangement, without the overhead of negotiation, vetting, or relationship management. A score-10 clipper automatically earns near-ceiling on every campaign.

No minimum follower requirements. The Trust Score already encodes a clipper's track record. A clipper with a high score and 500 followers who consistently generates 100K+ views per clip is more valuable than a clipper with 50K followers and no track record. The score reflects performance, not vanity metrics.

6.3 How the Trust Score Prevents Botting

A clipper who commits fraud by botting views and gets caught receives a DebtRecord and reputation slash. They can create a new account tomorrow, but the new Whop account and associated wallet starts at Trust Score 1, earning floor CPM. If the brand's range is \$0.50–\$3.00, the view botter just went from earning \$2.50/1K (score 8) to \$0.50/1K (score 1), an 80% pay cut. Rebuilding from score 1 to score 8 takes significant time, results, and consistency, where it's actually just more profitable to be honest.

The system closes both Sybil attack botting vectors: - Existing accounts are deterred by reputation loss. - New accounts are deterred by economics.

No KYC required for clippers. The trust score record IS the identity, harder to forge than an AI generated passport, because it requires months of consistent valued activity to build and seconds of fraud to destroy.

7 The Two-Dimensional CPM Model

Once a submission passes the confidence floor and is eligible, the clipper's effective CPM rate is determined by two independent factors:

Factor 1: Content Rewards Account Trust Score (1–10). The clipper's Content Rewards reputation, computed from account age, total verified views delivered across all campaigns, approval rate, fraud history, and consistency

of activity. A score of 1 represents a new account with no history or bad reputation. A score of 10 represents a top-tier clipper with years of activity and tens of millions of views. The score is dynamic: it drops immediately on activities that destroy the platform effectiveness, and recovers slowly over months of behavior that strengthens the system.

Factor 2: Confidence quality. How far above the floor the submission landed. A submission at exactly 80% earns the lowest CPM for that clipper's Trust Score. A submission at 95% earns significantly more. The quality factor measures the margin above the floor:

$$q = \frac{\text{confidence} - \text{confidence_floor}}{100 - \text{confidence_floor}}$$

The brand sets a floor CPM and ceiling CPM for the campaign, and selects a CR Curve shape: linear, convex (concentrates premium on top-tier clippers), or concave (makes the campaign more accessible to newer clippers). The clipper's effective CPM is:

$$r = \text{trust_score position on curve} \quad (0.0 \text{ to } 1.0, \text{ based on curve shape})$$

$$\text{combined_factor} = (r \times 0.50) + (q \times 0.50)$$

$$\text{effective_CPM} = \text{floor_CPM} + (\text{combined_factor} \times (\text{ceiling_CPM} - \text{floor_CPM}))$$

A clipper's effective CPM is driven equally by their Trust Score (who they are, 50%) and their submission's origin confidence (how good this specific clip is, 50%). Higher trust and higher fidelity both independently increase the rate.

7.1 Anti-Bot CPM Floor and Campaign Hijack Prevention

The confidence formula, the 80% floor, and the slider weights work together to prevent two categories of abuse:

View botting: The brand's CPM floor rate is calibrated by Content Rewards to be below the estimated cost of generating fake views that pass the campaign's specific eligibility requirements.

- The cost to bot is calculated from requirement multipliers: geo-filtering to Tier 1 countries, requiring brand name in handle, content uniqueness enforcement, and the origin verification itself all increase the cost of producing a fake qualifying submission.
- A volume discount factor adjusts the bot cost estimate downward for larger campaign budgets (bot services offer bulk pricing).
- The floor is set at 60% of the adjusted bot cost, ensuring view botting is a net loss at every scale point.

Campaign hijacking: - Consider a podcast with audio slider 80, visual slider 20: a meme visual hijack (95% audio, 3% visual) scores $(95 \times 0.80) + (3 \times 0.20) = 76.6\%$, below 80%, rejected. - Consider lifestyle content with audio slider 20, visual slider 85: a trending audio hijack (10% audio, 92% visual) scores $(10 \times 0.19) + (92 \times 0.81) = 76.4\%$, below 80%, rejected.

In both cases, whichever channel the brand cares about is the channel the hijacker can't fake. The formula mathematically prevents hijacking by weighting the confidence toward the channel the brand values.

Edge case protection: Content Rewards enforces a minimum combined slider value (audio + visual must be at least 80), to prevent situations where neither medium is weighted strongly enough to avoid hijacking.

Value-add from clippers is rewarded, not punished. A clipper who adds a scroll-stopping hook intro, engaging B-roll, captions, transitions, and an outro will score lower on origin confidence than a raw clip, but they earn on a per-view basis and their creative additions often drive significantly more views. The CPM model prices the tradeoff transparently: lower confidence = lower rate per view, but higher virality = more views. A B-roll heavy

podcast clip at 82% confidence on a score-7 account might earn \$1.80/1K views versus \$2.60 for a raw clip at 99%, but if the B-roll edit generates 2x the views, the clipper earns more total. The market finds the equilibrium between fidelity and virality, and optimizes accordingly.

● QUALITY FACTOR

$$Q = \frac{C - C_{\text{floor}}}{100 - C_{\text{floor}}}$$

Figure 10: Quality Factor: distance above the 80% confidence floor, normalized to [0, 1].



Figure 11: Content Rewards Curve Shape: how the chosen curve distributes CPM rewards across Trust Score levels.

● EFFECTIVE CPM

$$R_{\text{trust}} = f_{\text{curve}}(\text{Trust Score}) \in [0.0, 1.0]$$

$$F_{\text{combined}} = (R_{\text{trust}} \times 0.50) + (Q \times 0.50)$$

$$\text{CPM}_{\text{eff}} = \text{CPM}_{\text{floor}} + F_{\text{combined}} \times (\text{CPM}_{\text{ceil}} - \text{CPM}_{\text{floor}})$$

e.g. CPM floor = \$1 · CPM ceil = \$3 · brand selects curve shape

Figure 12: Effective CPM formula: floor CPM plus combined factor times the CPM range.

8 Submission Similarity and Account Scoring

When a clipper submits a clip, the system compares its fingerprint not only against the campaign's source material to verify origin, but also against every prior submission in that campaign to measure similarity. This second comparison layer detects when a clipper has reproduced another clipper's work rather than creating original content from the source material. The comparison uses the same perceptual fingerprinting infrastructure already in place (audio chromaprint matching, visual frame hashing, and temporal structure analysis) applied submission-to-submission instead of submission-to-source.

Rather than rejecting submissions that exceed a similarity threshold, the system feeds similarity signals into the clipper's account-level reputation score:

- A single instance of high similarity produces a minor score adjustment; it could be coincidental convergence.
- Repeated instances compound. A clipper who consistently submits content registering at 95%+ similarity to prior submissions, particularly when their Shelby read log shows no independent source access, accumulates a score profile that reflects systematic copying.
- The score degrades gradually, and with it, the clipper's CPM rate across all campaigns, not just the one where the similarity was detected.

This approach avoids the problems of hard rejection. Binary gates create adversarial disputes over where the threshold falls and train clippers to make the minimum modification necessary to pass rather than discouraging the behavior. A score-based system requires no bright line. A submission at 97% similarity adds more negative weight than one at 85%, which adds more than one at 70%.

Score recovery: A clipper who stops copying and produces original work with clean provenance chains (independent Shelby source access, reasonable editing timelines, meaningfully differentiated fingerprints) sees their score improve over time as recent positive history outweighs older similarity flags. The system rewards behavioral correction rather than applying permanent punishment.

Market fatigue prevention: By economically deprioritizing clippers who produce derivative content based on other clippers rather than the source material, the system naturally shifts payout weight toward original creators. The composition of active clippers on any campaign trends toward those producing diverse, differentiated content, not because anyone was banned, but because copying becomes less profitable than creating original content.

9 From Manual Gatekeeper to Automated Campaign Operator

9.1 Why Manual Arbitration Is the Bottleneck

Content Rewards' current system requires creators to manually review every submission: watching clips, checking URLs, verifying handles, counting views by country, judging content quality. A creator receiving 200 submissions

spends 10+ hours on manual review, inflating costs with virtual assistants who make inconsistent decisions about subjective requirements. The creator's time and money goes to back-office disputes and payroll admin costs, instead of creating more quality content that clippers can use to go viral.

Underneath this is a structural incentive misalignment: the creator is both the rule-setter and the judge. They define eligibility requirements (often vague), then decide who qualifies (often subjective). A creator can set "must be high quality" and then selectively reject compliant submissions using vague requirements as justification. The clipper can easily be scammed with limited recourse of issuing a costly chargeback.

When disputes arise, they move from the creator's inbox, to Content Rewards inbox, to Whop's inbox. At each step, a human with imperfect information or power makes a judgment call, or even worse, they are sent in a loop without a proper adjudication that destroys a user's trust, making it much more likely they will churn. More campaigns means more disputes means more support staff means higher costs. The dispute infrastructure scales linearly with volume, and at scale it breaks from operational systemic stress.

9.2 Why Shelby and Aptos Fix This

Shelby provides what centralized infrastructure cannot: neutral, tamper-proof evidence that no party controls. Content is stored on a decentralized network with cryptographic commitments on Aptos. Timestamps are in blocks confirmed by thousands of validators. Access control is in smart contracts, not admin panels. No party (not Whop, not Content Rewards, not the brand, not the clipper) can alter the evidence record after the fact.

Aptos provides a scalable settlement and coordination layer:

- Sub-second finality for streaming USDT payments to transfer payouts to clippers in real time
- The cheapest gas for hundreds of millions of monthly transactions, with a gasless user experience
- Campaign terms protected by a dual-speed change system: clipper-favorable updates take effect immediately, while clipper-adverse changes (lowering CPMs, tightening eligibility, removing source material) are delayed 24 hours so clippers always have advance warning before conditions change
- Safe smart contract design with fast eligibility checks and payout calculation based on perfect math

Together, Shelby and Aptos eliminate the human bottleneck that would otherwise scale linearly with dispute volume. The creator's role shifts from judge to campaign designer: set clear requirements, configure the CPM curve, deposit the budget, monitor results, and pursue campaign optimizations. The system handles everything else. 98% of submissions are processed with zero human input. The remaining 2% of genuinely ambiguous cases enter ADR, outsourced to market participants instead of support admin.

10 Real-Time Clipper Eligibility and Payouts

10.1 The Problem With the Old System

The current Content Rewards model requires every clip to be submitted for manual review — a human watches the clip, checks URLs, verifies handles, counts views, and makes a subjective eligibility judgment. This process is inconsistent and takes a few hours or days. A clipper who submits on Monday might not be made eligible until mid week, creating uncertainty that undermines platform trust.

The problem compounds when budgets run low and during viral spikes. The review queue becomes a lottery — two clippers submit equally eligible clips on the same day, one gets reviewed and paid, the other sits in queue until the budget is gone. During surges, the clippers generating the most views are stuck waiting for review while their viral window passes. The infrastructure breaks at exactly the moment Content Rewards is working as designed, denying payment to the clippers creating the most value.

A clipper who gets burned once — does the work, generates the views, and doesn't get paid because of queue timing or budget exhaustion during review — doesn't come back. And they tell other clippers. The payout

unpredictability, not the CPMs or the campaigns, is the core reason clipper talent churns. The trust layer breaks, and the flywheel stalls.

10.2 How Programmatic Verification Solves This

With Shelby and Aptos, the review queue doesn't exist. The clipper submits, the system verifies programmatically through dual-channel fingerprinting and confidence scoring, and USDT streaming begins immediately for every eligible view. No human intervention, no wait. The first clipper to submit and the hundredth are verified and earning at the same speed.

The remaining campaign budget updates in real time and is visible to everyone — clippers can see exactly how much is left before starting a new edit. When the budget approaches exhaustion, the creator is notified 24 hours before estimated depletion. They can top up to continue or let it exhaust naturally. Either way, every clipper was paid for every eligible view during the funded period. No one was denied payment because of the queue position.

In the old system, the question was "will I get reviewed before the budget runs out?" In the new system, the question doesn't exist. You submit, you're verified, you earn — every view, every time, in real time. The budget tail risk that plagued the manual system is eliminated by making eligibility deterministic and payouts instantaneous.

11 Account Farming and Acceptable Fraud Tolerance

11.1 The Attack

The most sophisticated fraud vector is farming accounts through months of legitimate participation to build a good reputation, then slowly supplementing organic views with undetectable levels of botted traffic. This is nearly impossible to fully prevent because the farming phase is indistinguishable from genuine clipping. The system does not attempt to eliminate this. Instead, it constrains the attack's profitability at every stage:

- The farming phase is expensive in time and creative effort.
- The siphoning phase is margin-constrained by detection thresholds that cap botted volume at levels too low to be highly profitable.
- The cumulative nature of the scoring system means even small repeated deviations gradually accumulate risk.

11.2 Account Reselling Detection

A secondary vector extends account farming further: farmed accounts with high reputation scores can be sold to operators who specialize in the siphoning phase, separating the labor of reputation building from the exploitation of that reputation. To mitigate account reselling, the system monitors for signals consistent with ownership transfer:

- Sudden changes in login IP, device fingerprint, browser cookies
- Account credential changes with Aptos private key rotation within Whop account
- Geographic access patterns or editing behavior that deviates sharply from the account's established baseline

These signals mirror the detection approaches used by social media platforms on Meta like Facebook, which flag account compromises and sales through the same environmental consistency checks.

11.3 Graduated Scoring Over Bans

Rather than suspending or banning accounts that trigger reselling signals, the system applies the same graduated scoring approach used across every other fraud vector: reselling indicators produce score reductions proportional to the severity and volume of the flagging.

- A single environmental change (a clipper logging in from a new IP while traveling, or switching to a new phone) produces a negligible adjustment that an honest account absorbs without meaningful impact to their CPM.
- Multiple simultaneous changes (new IP, new device, new browser cookies, credential update, Aptos wallet private key rotation, and a sudden shift in editing style all within the same week) produce a steeper reduction that reflects a much stronger likelihood of an ownership transfer.

This avoids creating a cliff edge where a legitimate clipper loses everything they have built because they went on vacation or upgraded their laptop, while still making account reselling economically unattractive because the buyer receives an account whose score has already degraded, reducing the CPM rates for campaigns below botting costs in response to the transfer signals, eroding the very reputation premium they paid for. Suspension and ban remain available under the terms of service as enforcement tools for cases where the evidence is overwhelming, but they are the final escalation, not the default response.

11.4 The Acceptable Outcome Is ROI Focused

The realistic steady state is a platform where the vast majority of activity is genuine, a small number of sophisticated operators extract a marginal amount, and the total fraudulent drain on campaign budgets stays in the low single-digit percentages, comparable to fraud rates tolerated by credit card networks and programmatic advertising. The true success metric is not fraud detection rate but creator retention: a campaign creator who sees strong returns and continues scaling their Whop campaign budgets is a creator for whom the system is working, regardless of whether every view was organic.

12 Phase 1: Objective-Only Requirements

12.1 What Phase 1 Means

At launch, every campaign requirement must be something Shelby and Aptos can verify programmatically with 100% determinism. No subjective requirements. No manual review. No approval queues. Every submission that meets the on-chain eligibility criteria is approved automatically and USDT streaming begins immediately. For example the creator name in bio, description URL, video length, tier-one views.

Every check is binary pass/fail. Every check is deterministic. The approval rate for eligible submissions is 100%. This is the strongest possible signal to attract clippers: you will get paid, every time, no exceptions, streaming USDT in real time.

12.2 Why Objective-Only at Launch

Phase 1 prioritizes clipper acquisition through absolute payout certainty. The current Content Rewards system's uncertainty (30-day holds, subjective rejections, creator rugging) suppresses participation. Removing all of it at once makes the platform the obvious choice for every clipper. More clippers means more submissions, more views generated, more USDT streaming, and more Content Rewards fees collected, which fills the insurance infrastructure needed for Phase 2.

Phase 1 also minimizes the dispute surface to as low as possible. If every requirement is objectively verifiable, there is nothing to dispute. A submission either passed the on-chain checks or it didn't. The smart contract is the authority.

The 80% confidence floor is specifically chosen to support this Phase 1 posture. It is strict enough to reject the overwhelming majority of fraudulent or off-topic submissions programmatically, while permissive enough to approve legitimate clips that include clipper value-adds (hooks, B-roll, captions, transitions).

13 Campaign Creation Fee Flow

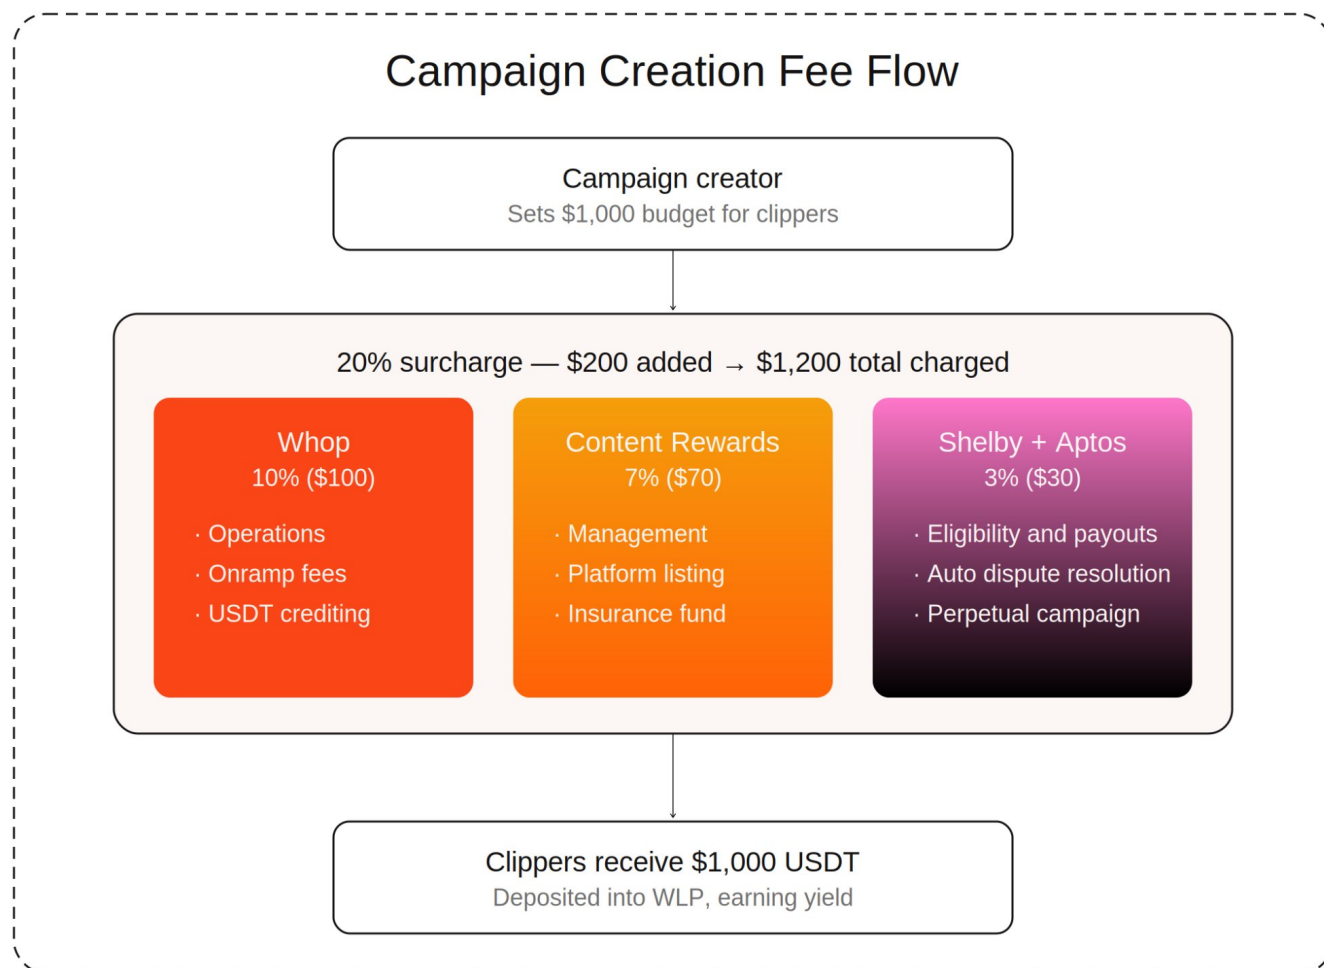


Figure 13: Campaign Creation Fee Flow: 20% surcharge split between Whop (10%), Content Rewards (7%), and Aptos (3%). Clippers receive 100% of the stated budget.

13.1 The 20% Surcharge

When a campaign creator sets a budget for clippers, a 20% surcharge is added on top. The creator sets what clippers will earn, and that amount is never touched. The surcharge funds the three layers of infrastructure that make Content Rewards work: Whop’s platform operations, Content Rewards management and insurance, with Shelby and Aptos eligibility, payments, plus a perpetual campaign that grows alongside to continue bringing in new creators and clippers.

Once the clipper submits content and qualifies as eligible, the USDT streaming payments start, where every view generates more money in real time, deposited into WLP vault earning yield.

13.2 Why the Creator Pays, Not the Clipper

The surcharge is additively charged to the creator on top of their chosen budget; clippers receive 100% of the stated amount.

- Clippers earn exactly what they see: A \$2.50 CPM means \$2.50 per thousand eligible views with no hidden deductions.
- Creator budgets are honest: A \$10,000 budget means \$10,000 flows to clippers; the creator knows their total

cost upfront (\$12,000).

- The surcharge funds the infrastructure that protects both sides: Whop operations, Content Rewards insurance, and Shelby + Aptos verification.

The creator pays for the system that makes their campaign deliver consistent ROI; and the clipper earns on a results basis per eligible view.

13.3 Whop: 10% Fee

Whop receives 10% of the creator’s stated budget (\$100 on a \$1,000 campaign). This funds:

- Operations: Hiring and training, growth initiatives, marketing spend, hosting and development that keeps Whop running smoothly in the background for everyone.
- Onramp fees: The raw costs associated with converting ACH, wires, or credit/debit card payments into USDT, while maintaining a safety net that accounts for fraud.
- USDT crediting: Upon payment confirmation from the creator, credit the campaign instantly by minting USDT directly from Tether at a 1:1 ratio with US dollars.

13.4 Content Rewards: 7% Fee

Content Rewards receives 7% of the creator’s stated budget (\$70 on a \$1,000 campaign). This funds:

- Management: Campaign operations, creator and clipper onboarding, documentation, and the product development that improves Content Rewards over time.
- Platform listing: The discovery infrastructure that connects creators with clippers such as campaign listings, search, recommendations, and the matching layer that ensures the right clippers find the right campaigns.
- Insurance fund (CRIF): The Content Rewards Insurance Fund that backstops dispute payouts and makes the entire system trustworthy enough for creators to scale budgets aggressively, while offering instant payouts to eligible creators every time, all the time.

13.4.1 Dynamic Utilization Curve

The 7% Content Rewards allocation is split internally between operating expenses and the CRIF. The share directed to the CRIF floats based on utilization, which is how much of the fund’s capacity is currently spoken for by active campaign exposure relative to its balance.

$$\text{utilization_ratio} = \frac{\text{total_active_campaign_exposure} - \text{CRIF_balance}}{\text{total_active_campaign_exposure}}$$

- When the CRIF is healthy (low utilization), the CRIF takes a smaller share of the 7% fee total, at approximately 1.5%. Content Rewards retains more operating margin. Excess overflows to the WLP as depositor yield.
- When the CRIF is stressed (high utilization), the CRIF takes a larger share, of up to 4% or more. Content Rewards earns less margin temporarily, but the reserve rebuilds faster. Higher premium flow to the WLP attracts external capital through elevated yield.

The curve self-balances. The creator never sees or feels the internal split. They always pay the same 20% surcharge. But the system underneath breathes with market conditions, automatically strengthening the insurance infrastructure when it needs it most and releasing capacity to operations when the fund is healthy.

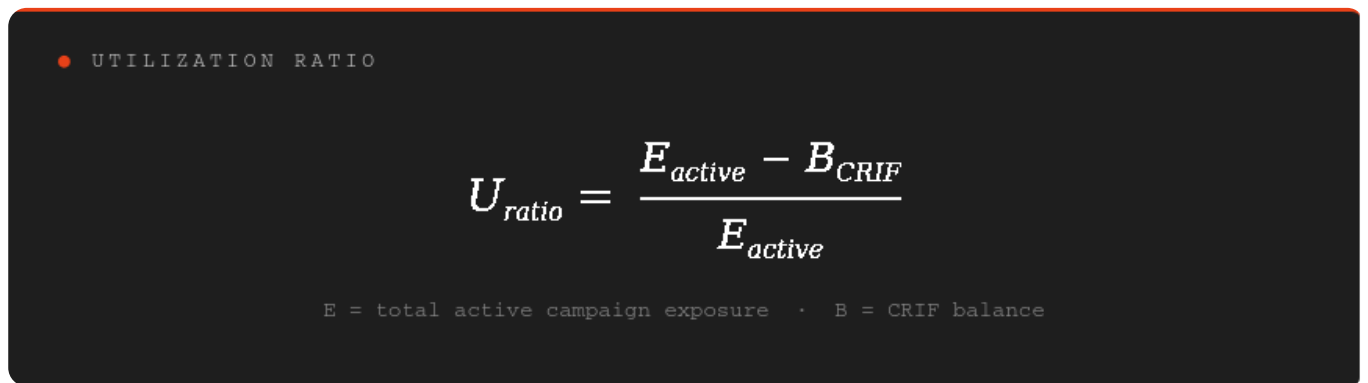


Figure 14: Utilization Ratio: CRIF allocation self-balances between insurance reserve and WLP yield based on fill level.

Three Utilization States

Healthy ($u < 0.3$) — CRIF allocation $\approx 1.5\%$ of the 7% fee. WLP overflow $\approx 5.5\%$. The insurance reserve is well-capitalized. Content Rewards retains maximum operating margin. Depositors earn steady overflow yield.

Phase 1 Target ($u \approx 0.55$) — CRIF allocation = 2.8%. WLP overflow $\approx 4.2\%$. The system is at its designed equilibrium — insurance builds at a sustainable rate, operations are funded, and yield flows to depositors.

Stressed ($u > 0.8$) — CRIF allocation $\approx 4\%+$. WLP overflow $\leq 3\%$. A dispute spike or large claim has drawn down the reserve. The curve automatically pulls more premium into the CRIF. Elevated yield attracts external capital to the WLP.

13.4.2 CRIF Overflow to WLP

When the CRIF reaches its target balance (the trailing 7-day average claim volume), the excess premium revenue overflows into the WLP as yield for depositors. This overflow is the mechanism that makes the WLP attractive to external capital where depositors earn insurance premium yield that reflects real-time system health. A 7-day rolling average means the CRIF target recalibrates at the same cadence as the ADR flow itself; a spike in disputes this week immediately raises the target next week, and the utilization curve responds by pulling more premium into the fund.

13.5 Aptos: 3% Fee

Aptos receives 3% of the creator’s stated budget (\$30 on a \$1,000 campaign).

- Eligibility and payouts: Shelby verification stack for fingerprinting, confidence scoring, and programmatic eligibility; Aptos settlement layer for streaming USDT via micropayment channels with sub-second finality.
- Auto dispute resolution: On-chain smart contracts for dispute filing, lien placement, evidence storage, voting, and resolution which are all gas-sponsored for all users.
- Perpetual campaign: Self-funding Aptos × Whop × Content Rewards campaign paying clippers for generating platform content, paid in APT for believer natural self-selection.
- Onchain purchases: The 3% is converted from USDT to APT through an onchain CLOB or PropAMM. No off-chain exchanges, no CEX custody, no opaque OTC desks. Every conversion is publicly verifiable on the Aptos blockchain.

The table below consolidates every dollar on a representative \$1,000 creator budget — the same numbers referenced individually in the three subsections above.

Recipient	Rate	Amount	What It Funds
Whop	10%	\$100	Platform ops, onramp processing, USDT minting at 1:1 with Tether
Content Rewards	7%	\$70	Campaign management, platform listing, CRIF insurance reserve
Aptos	3%	\$30	Shelby verification, ADR smart contracts, Perpetual Campaign
Clippers	—	\$1,000	Per-view streaming payments, zero deductions from stated budget
Creator total cost		\$1,200	\$1,000 budget + \$200 surcharge (20%)

Table 2: Fee breakdown on a \$1,000 creator budget. The 20% surcharge is additive and paid entirely by the creator. Clippers receive 100% of the stated amount with zero deductions.

14 Perpetual Campaign

The Perpetual Campaign is a Content Rewards campaign that never ends, and never runs out of budget. Run by Aptos as the creator, making it the most reliable campaign where the platform markets itself through its own product, funded by its own fees, without a new marketing budget.

The source material is Whop and Content Rewards itself with how the platform works, how clippers earn, what their results look like, and how it’s being powered by Aptos. Every clip that goes viral is a direct acquisition channel bringing new clippers and creators into the ecosystem.

14.1 Dynamic APT Budgeting

When the platform is quiet and needs growth most, the dynamic allocation shifts more funding toward the perpetual campaign. When the platform is surging, it scales back because organic growth is already happening, and more gas fees are being spent on streaming payments.

High Activity (Deficit): When Content Rewards usage is surging, gas draining fast from all the sponsored settlements and ADR transactions. Treasury routes up to 80% to gas, Engine scales back to 20% but never stops, with always a live campaign for new clippers to join.

Low Activity (Surplus): When Content Rewards usage is low, the gas accounts are well-funded. Treasury routes up to 80% to The Engine, maximizing growth budget when it matters most. The Perpetual Campaign attracts new clippers, content is viral, activity picks up, gas demand rises, Treasury naturally rebalances back, priming the next cycle without anyone making a decision.

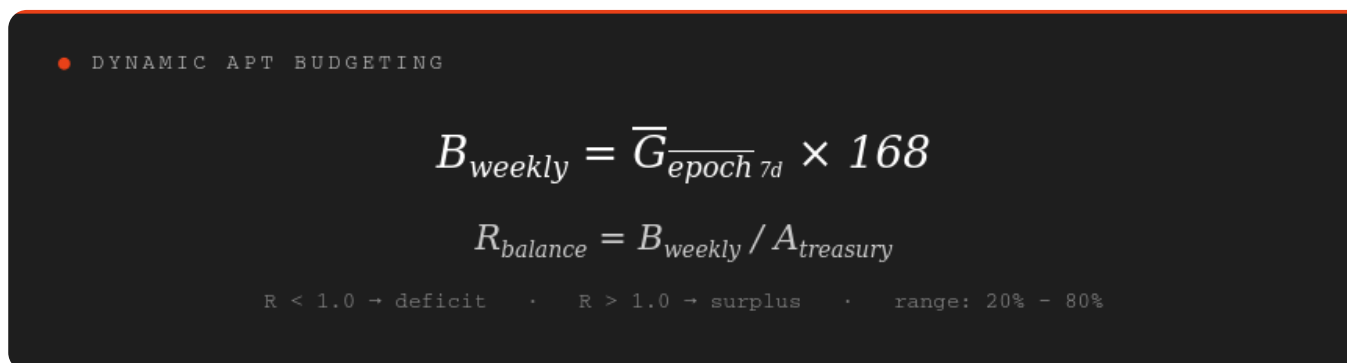


Figure 15: Dynamic APT Budgeting: the 3% Aptos fee self-balances between gas sponsorship and the Perpetual Campaign based on platform utilization.

15 Creator Trust Score

15.1 What the Creator Trust Score Encodes

The Creator Trust Score (1–10) encodes whether a creator is someone clippers and the platform should want to work with. The score is entirely within the creator’s control — it reflects only decisions the creator made, never external factors like clipper behavior, CRIF capacity, or market conditions.

The score captures three dimensions of creator quality: whether they honor the economic contracts they offer (reliability), how they treat clippers through campaign management decisions (partner quality), and whether their investment in the platform is growing or shrinking over time (commitment trajectory).

15.2 The Three Dimensions

Dimension 1: Contract Reliability. Does the creator follow through on what they promised? This is measured through campaign completion rate (what percentage of budgets reach natural exhaustion versus premature closure), dispute filing impact (do their disputes catch real fraud or drain the system), and adherence to the terms they set. A creator who launches a campaign, attracts 200 clippers, and lets the budget run to exhaustion scores very differently from one who launches, attracts 200 clippers, and closes with 60% remaining.

Dimension 2: Partner Quality. Is this creator good to work with, or do they create a hostile environment for clippers? This is measured through the frequency and severity of negative campaign adjustments, how often the creator raises confidence floors, lowers CPM ranges, removes source content, or tightens eligibility criteria in ways that disadvantage clippers who are actively working on submissions. A creator who launches a well-configured campaign and lets it run with minimal changes scores high. A creator who launches and then spends the next two weeks constantly tightening parameters, raising the floor by 3 points here, removing a source video there, dropping the CPM ceiling by \$0.20, is creating an unstable environment where clippers cannot plan their work. Even if every individual change respects the delay window, the cumulative churn degrades the score. The system tracks a negative adjustment frequency metric: how many clipper-adverse changes per week, weighted by severity, normalized against the campaign’s active duration.

Dimension 3: Commitment Trajectory. Is this creator scaling up or winding down? This is measured through budget growth over time, a trailing comparison of the creator’s campaign spending across rolling 30/60/90-day windows. A creator whose average campaign budget grew from \$2K to \$8K to \$25K over six months is demonstrating that the platform works for them and they are reinvesting their returns. A creator whose budgets are shrinking from \$10K to \$5K to \$1K is signaling declining confidence. The score rewards sustained growth, is neutral on flat spending, and penalizes sustained decline. Single outlier campaigns (one large test, one small experiment) are smoothed by the rolling window; the signal is the trend, not any individual campaign.

This dimension matters because a growing creator is a healthier counterparty for everyone in the system. Clippers prefer creators who are scaling (more budget means more earning opportunity and longer campaign life). Delegates prefer creators who are scaling (more fee revenue from larger campaigns). The CRIF prefers creators who are scaling (more premium income from growing campaign exposure). And Content Rewards prefers creators who are scaling (more fee revenue). The score should reflect that alignment.

16 Clipper Trust Score

16.1 What the Clipper Trust Score Encodes

The Clipper Trust Score (1–10) encodes whether a clipper produces legitimate content and delivers genuine value to campaigns. A score of 1 represents a brand-new account with no history or a compromised reputation. A score of 10 represents a top-tier clipper with sustained activity, tens of millions of verified views, and a clean track record. The score is dynamic: it drops immediately on fraud, copying, or account anomalies, and recovers slowly over months of consistent legitimate work.

16.2 What Feeds the Score

The score is computed from multiple signals that compound over a clipper's full history:

- Account age and consistency: Longer track records with steady activity carry more weight than short bursts, because sustained behavior is harder to fake.
- Total verified views delivered: Cumulative views across all campaigns, verified through Shelby's read-metering, not self-reported platform metrics.
- Approval rate and confidence quality: The percentage of submissions that pass the confidence floor, and how far above the floor they land on average. A clipper who consistently submits at 90%+ confidence signals higher fidelity than one who hovers at 81%.
- Submission similarity: Repeated high-similarity submissions to prior clips in the same campaign degrade the score gradually, as described in the Submission Similarity section. Original work rebuilds it.
- Fraud history: Any DebtRecord, lien seizure, or guilty ADR verdict produces an immediate and severe score reduction. Recovery requires extended periods of clean activity.
- Account environment consistency: Sudden changes in IP, device, credentials, or editing behavior consistent with account reselling produce graduated score reductions proportional to the severity of the signals.
- WLP deposit reputation: Sustained deposits in the WLP contribute positively over time, reflecting long-term commitment to the platform's health.

16.3 How the Score Affects Earnings

The Clipper Trust Score directly determines the clipper's position on the campaign's CPM curve. A score-1 clipper earns floor CPM on every campaign. A score-10 clipper earns near-ceiling. This is the mechanism that replaces gatekeeping with pricing — every campaign is open to every clipper, but the economic reward scales with proven reliability.

17 Campaign Trust Score

17.1 What the Campaign Trust Score Encodes

The creator Trust Score is a personal credit rating — it travels with the creator across all their campaigns. The Campaign Trust Score is a real-time health meter for a specific active campaign, visible to all participants. It answers a different question: if something goes wrong on this campaign right now, can the system handle it?

The Campaign Trust Score recalculates continuously as the campaign progresses. It is not under any single party's control — it is a composite of creator behavior, system capacity, and clipper behavior on this specific campaign.

17.2 Campaign Trust Score Inputs

Creator Trust Score feeds in as the behavioral floor. A high-trust creator sets a higher baseline for all their campaigns. A low-trust creator drags the starting point down. This is the only input that is not campaign-specific.

Depth factor measures this campaign's budget relative to CRIF available capacity. A \$5K campaign against a \$500K CRIF barely registers. A \$200K campaign against the same CRIF represents significant concentration risk. This ensures the score captures systemic risk even for well-behaved creators running outsized campaigns. The formula:

$$\text{depth_factor} = \frac{\text{campaign_budget}}{\text{CRIF_available_capacity}}$$

Solvency ratio tracks what percentage of clipper earnings from this specific campaign are still sitting in the WLP as lien-able backstop versus having been withdrawn. This is the most dynamic input. If disputes are resolved primarily through liens on guilty clippers' WLP positions with the CRIF as the backstop only if those clippers already withdrew, then the solvency ratio directly measures how recoverable a potential dispute would be. A solvency ratio of 0.85 means 85% of payouts are still in the WLP and can be lien-ed. A solvency ratio of 0.15 means 85% has been withdrawn and any dispute would fall primarily on the CRIF, potentially trickling into socialized WLP losses.

Clipper quality on this campaign captures the aggregate trust profile of active submitters. A campaign where the average clipper Trust Score is 7.5 and submissions show low similarity clustering is healthier than one where the average is 3.2 with high similarity clustering (suggesting coordinated bot-like behavior). This input is not controlled by the creator or any individual clipper; it reflects the emergent behavior of the campaign's participant pool.

● DEPTH FACTOR

$$D = \frac{\text{Budget}_{\text{campaign}}}{\text{Capacity}_{\text{CRIF}}}$$

Figure 16: Depth Factor: campaign budget relative to CRIF available capacity—the primary systemic risk input to the Campaign Trust Score.

18 Campaign Lifecycle Actions

Six distinct actions can be taken on a live campaign. The table below shows timing, financial outcomes, and Trust Score implications at a glance — the subsections that follow detail each action in full.

Action	Takes Effect	Creator Keeps	Clipper Outcome	Trust Score
Instant change	Immediately	Ongoing	Better terms, no delay	+ Builds over time
Delayed change	After 24 h	Ongoing	Original terms 24 h, then new	- Mild per change
Budget top-up	Immediately	Ongoing	Earnings window extended	+ Commitment
Sunset	7-day decay	90% minus haircut	CPM decay over 7 days	Haircut-dependent
End Campaign	Immediately	50% of remainder	50% pro-rata, zero notice	-- Steepest
Dispute	Immediate halt	ADR (≤ 4 wks)	Lien; released or seized	+ proven / -- empty

Table 3: Campaign lifecycle actions at a glance. Optimization actions (green/amber) keep the campaign running. Sunset (gray) winds down gracefully. End and Dispute (red) terminate immediately.

18.1 Campaign Optimization (Most Common)

All changes to a running campaign, such as CPM ranges, confidence floors, eligibility criteria, source material, audio/visual fidelity, all flow through a single optimization action with two speeds:

Instant changes (clipper-favorable): Any update that makes it easier for existing clippers takes effect immediately and positively impacts the creator Trust Score over time. This includes raising CPMs, removing eligibility criteria, lowering the confidence floor, editing audio/visual fidelity while keeping the ratio the same, and adding new source material to the campaign.

Delayed changes (clipper-adverse): Any update that makes it harder for existing clippers takes effect after a 24-hour delay and negatively impacts the creator Trust Score over time. During the delay, clippers continue earning under the original terms. This includes lowering CPMs, editing or adding eligibility criteria, raising the confidence floor, editing audio/visual fidelity with a different ratio, and removing existing source material from the campaign.

Mixed changes (simultaneously easier and harder): Delayed by default. If a creator raises CPMs but also adds new eligibility criteria in the same update, the adverse component governs the timing, and the entire update takes 24 hours.

The frequency of adverse changes is tracked in the creator Trust Score’s partner quality dimension. A creator who queues adverse changes every other day degrades their score even though each individual change respects the delay window. The cumulative churn is the signal.

Budget top-up: When a campaign’s budget is estimated to expire shortly based on trailing payout cadence, the creator is notified. No penalty for not topping up; instead the campaign exhausts naturally. Consistent top-ups with significant size contribute positively to commitment trajectory dimension and work toward fee reductions on the total 20% surcharge over time, scaled by cumulative campaign volume and creator Trust Score, reaching economies of scale.

18.2 Sunset Campaign (Graceful Wind-Down)

The creator initiates a 1-week sunset. The campaign is publicly marked as ending with the exact date visible to all clippers. During the sunset week, all eligible submissions — both prior submissions still accruing views and

new submissions made during the period — continue earning with linear CPM decay: day 1 pays at effectively the original CPM value, day 7 pays at effectively zero.

The sunset haircut determines the maximum amount the creator can lose from their remaining unspent budget. The haircut portion funds the sunset week's decaying CPM payouts. If clippers don't fully consume the haircut portion during the 7-day decay (which is likely — participation drops as CPMs approach zero), whatever remains also returns to the creator. The non-haircut portion returns to the creator immediately for future campaigns or withdrawal.

18.2.1 Sunset Haircut Formula

● VIRALITY FACTOR

$$V = \frac{1}{1 + e^{-k_v \times (r_v - 1.0)}}$$

r_v = current daily views / campaign average daily views · $k_v \approx 2.5$

Figure 17: Virality Factor: input to the sunset haircut formula.

● BUDGET REMAINING FACTOR

$$B = \frac{1}{1 + e^{-k_b \times (p_{rem} - 0.50)}}$$

p_{rem} = remaining budget / total campaign budget · $k_b \approx 8.0$

Figure 18: Budget Remaining Factor: more budget remaining means higher haircut exposure.

● TRUST SCORE FACTOR

$$T = \frac{1}{1 + e^{k_t \times (S_{creator} - 5.0)}}$$

inverted — higher trust = lower haircut · $k_t \approx 1.5$

Figure 19: Trust Score Factor: higher-trust creators pay a lower haircut percentage.

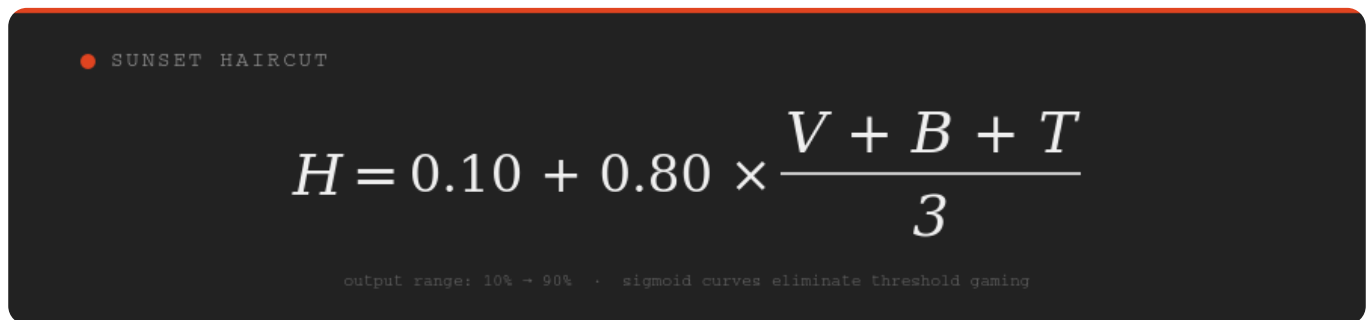


Figure 20: Sunset Haircut formula combining virality, budget remaining, and trust score factors.

Three factors, each normalized to 0→1 using sigmoid curves, combined into a single haircut output on the 10%→90% range. Sigmoids eliminate edge gaming — there are no thresholds to sit just above or below, and every marginal change in input produces a marginal change in output.

The table below shows how each factor behaves at its extremes and which direction it pushes the haircut *H*.

Factor	Low Input	High Input
F_v (Virality)	Campaign is quiet — few active clippers, low view velocity. <i>H stays near 10%</i> ; creator keeps more.	Campaign is viral — many clippers producing content, high views. <i>H rises toward 90%</i> ; compensates active clippers.
F_b (Budget remaining)	Most budget already spent — the campaign ran its course. <i>H stays low</i> .	Most budget still unspent — creator barely let it run. <i>H rises</i> ; penalizes early exit.
F_t (Trust, inverted)	Creator has a <i>high</i> Trust Score — reliable, consistent history. <i>H stays low</i> ; rewards good faith.	Creator has a <i>low</i> Trust Score — pattern of adverse actions. <i>H rises</i> ; penalizes bad actors.

Table 4: Sunset haircut factor behavior. Each factor is sigmoid-normalized to [0, 1] before combining. The inversion on F_t means higher trust *reduces* the haircut — reliable creators always pay less.

18.3 End Campaign (Immediate)

The campaign is immediately closed and disappears from public listing on Content Rewards. No sunset window, no decay period. This is for creators who want an instant, clean break.

The economics: the creator recoups 50% of the remaining unspent clipper budget, and 50% is redistributed pro-rata to clippers based on their share of total campaign earnings before closure, deposited into their WLP positions. The creator’s 50% is returned to their Whop account for future campaigns or withdrawal.

Instant closure carries the steepest Trust Score hit of any non-dispute action, because it gives clippers zero time to adjust. A creator who ends campaigns instantly and repeatedly will see their partner quality and contract reliability dimensions degrade rapidly.

The fixed 50/50 split creates natural sorting with the sunset option: a good-faith creator (sunset haircut no more than 10%) always prefers sunset (keeps up to 90%) over instant end (keeps only 50%). A predatory creator facing a sunset haircut up to 90% likely chooses an instant end (keeps 50%) because it’s less punishing, which is better for clippers because they get an immediate payout of 50% without an additional week of effort for low quality campaigns with limited future payouts.

Why the Sorting Works: \$5,000 Remaining Budget

Good-faith creator (haircut = 10%):

- Sunset: keeps up to \$4,500 (90%). Clippers get 7-day CPM decay funded by the \$500 haircut portion.
- End Campaign: keeps \$2,500 (50%). Clippers get \$2,500 pro-rata immediately.
- *Always prefers Sunset — saves \$2,000, and the Trust Score penalty is minimal.*

Predatory creator (haircut = 90%):

- Sunset: keeps only \$500 (10%). \$4,500 funds clipper decay.
- End Campaign: keeps \$2,500 (50%). Clippers get \$2,500 immediately.
- *Prefers End Campaign — but clippers still win: immediate \$2,500 with no wasted week.*

The system sorts creators into the action that is best for clippers without Content Rewards making the decision. Good-faith creators self-select into Sunset; bad-faith creators self-select into the option that still protects clipper earnings.

18.4 Dispute Campaign (Last Resort)

File a dispute when the creator genuinely believes they have been defrauded. Unlike all other actions, the dispute halt is immediate with no grace period, no reduced-rate streaming, all because the creator is asserting that ongoing payouts may be going to fraudulent activity.

The dispute does not halt all of the creator's campaigns. Only the disputed campaign is halted, plus any other campaigns with source material overlap above the duplicate detection threshold. Genuinely different campaigns continue running normally. This prevents creators from cloning a disputed campaign to keep running while ADR plays out, without punishing scaling creators who operate across multiple distinct content categories.

The dispute filing fee is calibrated relative to the campaign's total budget and not a flat dollar amount. Scaling the fee as a percentage of campaign budget (with a floor and ceiling) ensures the seriousness filter works at every scale, which also properly incentivizes the bounty hunter.

The creator's dispute filing impact metric in their Trust Score tracks filings aggressively: disputes where no bounty hunter claims the case, or that resolve with zero guilty verdicts, accumulate score damage that compounds with each empty filing. A creator with a genuine fraud claim is fully compensated through the ADR process, and confirmed claims paid from lien seizures and CRIF backstop, and the creator's Trust Score is positively impacted by successful filings.

19 Bounty Hunters

Bounty hunters are independent investigators who claim dispute cases and compile evidence for voter review. Shelby stores content verification data that isn't visible on-chain (read metrics, bot probability scores, view verification logs, payout history) and access to this data can be token-gated behind Shelby for bounty hunters specifically. The exact access mechanism is still in consideration.

19.1 Compensation Model

Guaranteed retainer: Funded by the flat dispute filing fee the creator pays upfront. This goes directly to the hunter who claims the case, regardless of outcome. Investigation takes real time, and without upfront compensation the hunter is working on spec, optimizing for speed over quality. The retainer ensures thorough investigation because the hunter is paid for the work itself.

Performance reward: A percentage of the clawed-back amount if WLP voters agree with the hunter's findings. This incentivizes getting it right.

Proportional staking: When a bounty hunter claims a case, they stake an amount proportional to the total lien value they are placing on the implicated clippers, scaling with the magnitude of disruption they are causing to other participants' earnings. If voters agree, the hunter earns the retainer plus the performance reward and their stake is returned. If voters disagree, the hunter keeps the retainer (they did the work) but forfeits their stake into the CRIF's Insurance Reserve and earns no performance reward.

Bounty Hunter Trust Score: Hunters build an on-chain reputation score based on their win/loss ratio weighted by claim size, with a precision penalty for over-flagging addresses. High-accuracy hunters attract the most cases because their evidence is trusted by voters. Low-accuracy hunters stop getting claims because their track record signals unreliability.

Career path: The ideal bounty hunter is someone who already understands clipping. A clipper earning consistent income on Content Rewards knows what legitimate engagement looks like because they see it daily. They can spot bot views, fake accounts, and inflated metrics faster than any outside analyst. This creates a natural career path within the ecosystem: clip, earn, then investigate.

Market mechanism: Bounty hunters replace the support queue with a bounty board. Cases get claimed because solving them pays. The harder the case, the larger the stake, the more attractive it is to investigate. Resolution speed is set by market incentive, not by how many support agents are on shift.

20 Automated Dispute Resolution (ADR)

ADR is the market-driven mechanism that replaces manual arbitration. The system operates on the principle that every clipper is innocent until proven guilty: they have already been paid because the automated verification system approved their submissions, and the burden falls on the bounty hunter to demonstrate that the approval was wrong.

A creator files a dispute when they identify potential fraud where they've overpaid based on campaign eligibility verification. The creator pays a dispute filing fee upfront designed as a seriousness filter. If the fee is \$200 and the dispute is over a \$150 overpayment, the creator doesn't file. This self-selects for disputes where the creator genuinely believes they've been materially defrauded of a significant amount relative to their ROI. After the filing, all creator's campaigns are paused during ongoing dispute, creating an additional operational deterrent against frivolous claims, where the opportunity cost is presented clearly, and minor inconveniences aren't clogging up ADR.

Core design principle: Each participant acts exactly once. No back-and-forth, no appeals queue, no support tickets. One action per participant, then the system auto-resolves. The entire ADR completes within 4 weeks maximum, ensuring genuine claims with only substantial losses are prioritized.

20.1 The ADR Flow and Timeline

- 1. Dispute Filing.** A creator files a dispute, identifying the campaign and the suspicious activity. The creator pays the filing fee. The creator does not need to specify which clipper addresses are responsible; that is the bounty hunter's job. A dispute is opened, shown only to bounty hunters, for a maximum period of up to 1 week, balancing paused campaign opportunity cost with adjudication.
- 2. Claim Case.** Bounty hunters evaluate the claim. If a hunter believes the claim has merit, they claim the case, stake proportionally to the anticipated lien value, and receive the filing fee as a guaranteed retainer. Full evidence access is unlocked via Shelby. If no hunter claims the case within the window, the filing fee is redirected to the CRIF and the dispute does not proceed. No hunter claiming the case means no one found sufficient evidence to investigate, or that the alleged losses were financially trivial to pursue further, not that the creator is right by default. Up to 1 week.
- 3. Investigate.** If claimed, the hunter builds the case: querying Shelby read logs, cross-referencing platform APIs, analyzing wallet patterns, and compiling findings. Critically, the hunter must identify the specific clipper wallet addresses whose submissions constitute the alleged fraud, along with the USDT value attributed to each address. Up to 1 week.
- 4. Submit and Lien.** The hunter uploads their evidence report to Shelby (immutable and timestamped) and submits the list of implicated clipper addresses. This submission executes as a single atomic transaction using Aptos script composer: in one on-chain action, the evidence is uploaded to Shelby, the implicated addresses are recorded in the dispute contract, and a proportional lien is placed on each implicated clipper's WLP position corresponding to their share of the alleged overpayment. Because the investigation happens off-chain and the lien drops in the same transaction as the evidence submission, the implicated clippers have no window to front-run.
- 5. Counter-Evidence.** Implicated clippers receive a bounded counter-evidence window before voting opens. Each clipper can review the bounty hunter's full evidence package and upload their own counter-evidence to Shelby. A clipper who does not submit counter-evidence within the window is deemed guilty by default. A clipper who submits counter-evidence has their case evaluated on the merits by voters. Up to 1 week.
- 6. Vote.** Voting opens immediately after the counter-evidence window closes. WLP depositors review both the bounty hunter's evidence and any counter-evidence, then vote on each implicated clipper individually: guilty or not guilty. Votes are weighted by a conviction staking model where depositors optionally lock Whop USDT behind their vote, with maximum weight determined by both deposit size and Content Rewards Trust Score. If no depositors vote at all, all clippers are deemed innocent and all liens are released. The system defaults to innocence in the absence of an active decision to convict. Up to 1 week.
- 7. Resolve.** The smart contract executes the outcome per-clipper: - **Guilty:** Lien converts to seizure. Funds transfer to the CRIF's Lien Escrow. The Lien Escrow disburses to the creator. - **Not guilty:** Lien releases in full. WLP position returns to normal. Trust Score unaffected.

The creator receives only the portion attributable to guilty clippers, never the portion attributed to cleared clippers. The CRIF's Insurance Reserve covers the gap only when a guilty clipper cannot pay, never when an innocent clipper is cleared. This prevents insurance fraud where a creator colludes with wallets to extract CRIF funds for non-fraudulent claims.
- 8. Finalize.** The dispute is permanently closed and logged as historical cases for bounty hunter training purposes, and the money is fully settled at the end of an Aptos epoch, up to 2 hours.

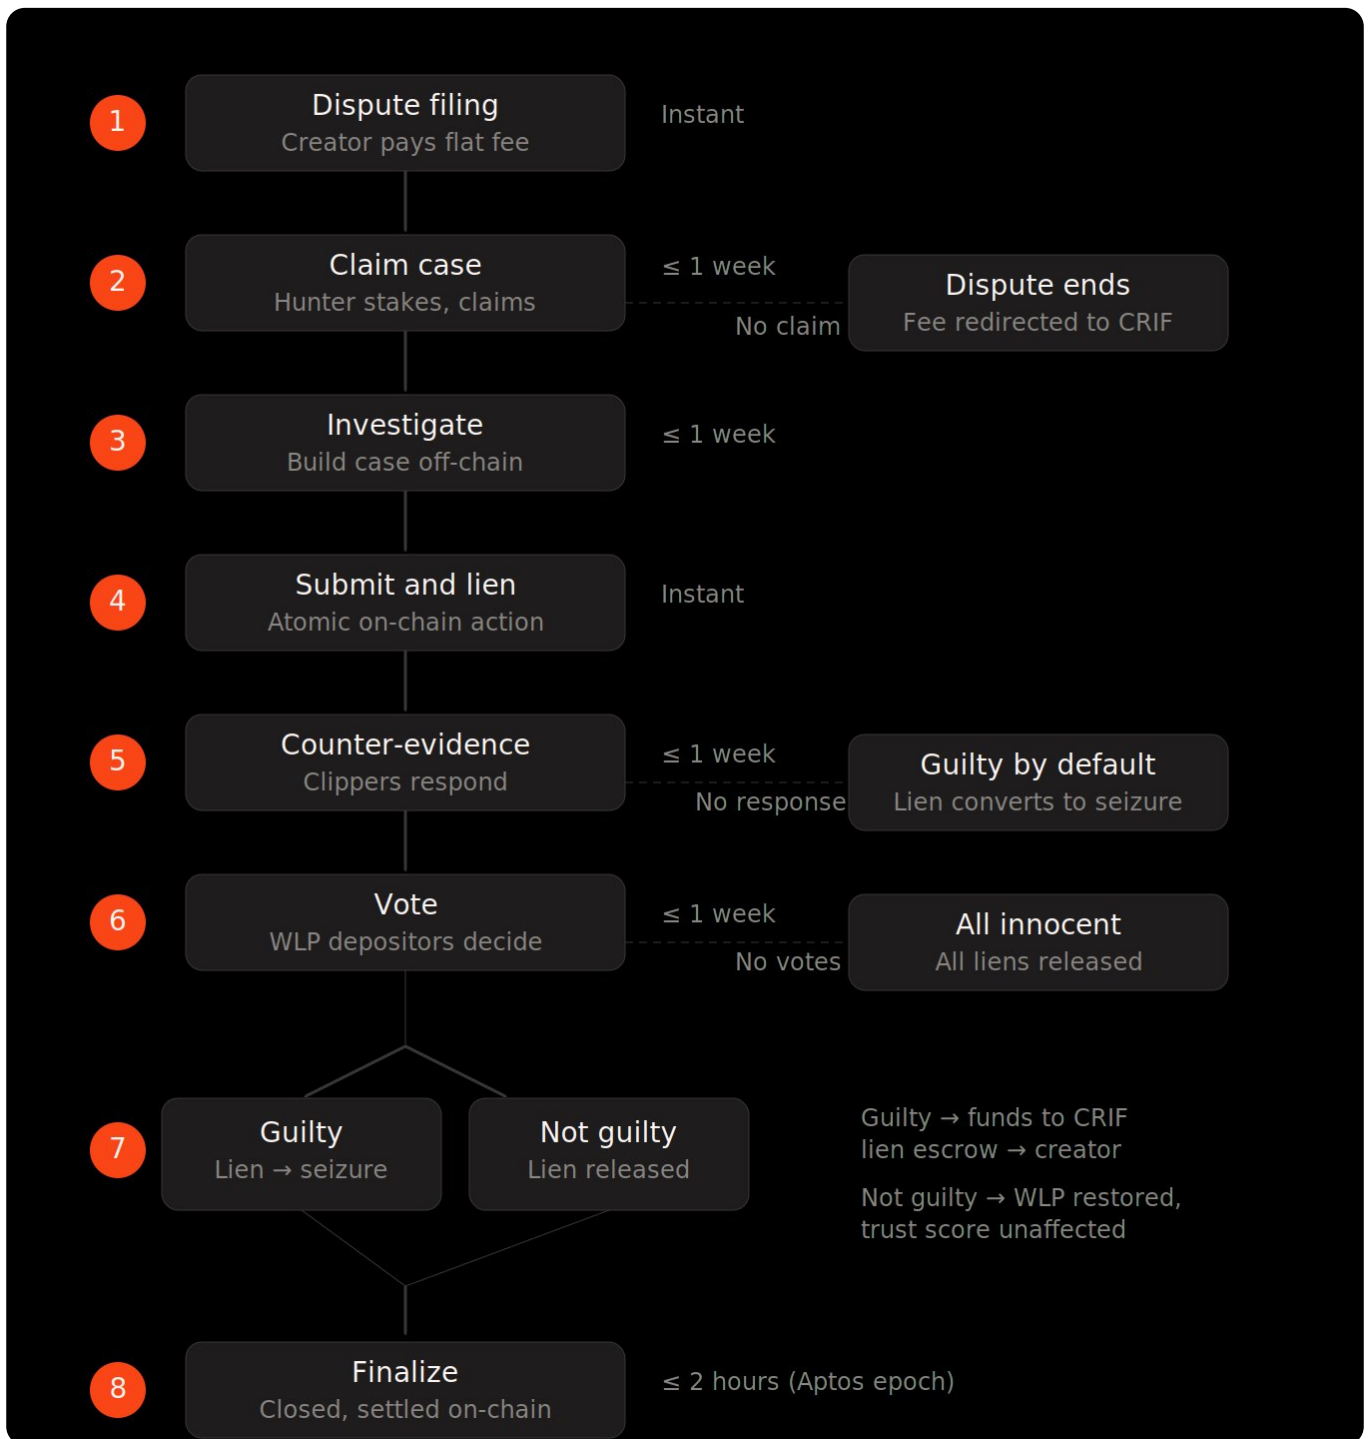


Figure 21: ADR Flow: from filing through investigation, lien, counter-evidence, vote, and resolution.

21 Content Rewards Insurance Fund (CRIF)

The Content Rewards Insurance Fund operates as two distinct on-chain sub-ledgers under a single contract, each with a clear purpose and separate accounting:

21.1 Insurance Reserve

- Funded by the Content Rewards fee allocation through the dynamic utilization curve described in the Campaign Creation Fee Flow section.

- This is the protocol's own capital: money that Content Rewards has earned and set aside to cover dispute costs.
- Target balance is calculated as the trailing 7-day average claim volume.
- While the reserve is below target, the utilization curve directs a larger share of the 7% Content Rewards fee toward the CRIF.
- Once it reaches the target, excess fee revenue overflows into the WLP as additional yield for depositors.
- The Insurance Reserve only covers the gap left by guilty clippers who cannot pay. It never covers the portion attributed to clippers voted not guilty.

21.2 Lien Escrow

- Holds seized clipper earnings from confirmed guilty verdicts, pending disbursement to the creator who filed the dispute.
- Funds enter the Lien Escrow when a guilty clipper's lien is converted into a seizure and exit when the dispute resolution finalizes and the creator is paid out.
- The Lien Escrow is a pass-through that never accumulates a permanent balance.

At any moment, anyone can query the CRIF on-chain and see the Insurance Reserve balance (protocol capital available for future disputes) separately from the Lien Escrow balance (clipper earnings in transit to creators for resolved disputes). The two pools are never commingled.

When an Insurance Reserve draw-down occurs, the utilization curve automatically redirects a larger share of the 7% fee to refilling the reserve. Normal overflow to the WLP resumes once the target is restored.

22 Whop Liquidity Provider (WLP)

The WLP is a USDT vault on Whop powered by Aptos Move smart contracts that anyone can deposit into. Depositors receive Whop USDT tokens representing their share. The vault earns yield from two sources: overflow premium from the CRIF (when the CRIF is at target balance, excess fee revenue flows to the WLP as yield) and participation in ADR voting.

22.1 Clipper Earnings and Automatic WLP Deposit

Every clipper's verified earnings flow directly into the WLP as their default settlement destination. Rather than streaming raw USDT on-chain for every view (which would produce thousands of gas-inefficient microtransactions during periods of steady viewership and awkward dust-sized settlements during plateaus), the system uses Shelby micropayment channels to aggregate earnings off-chain and settle into the WLP at regular intervals.

How the settlement works:

- As a clipper's eligible views accrue, Shelby's native micropayment metering tracks the exact USDT owed in real time within an off-chain payment channel between the campaign's escrow and the clipper. The clipper's dashboard reflects their accumulating balance as it grows; the experience feels like real-time streaming even though the on-chain settlement is batched.
- At regular intervals, Shelby closes the payment channel and settles the net balance into the WLP in a single Aptos transaction, crediting the clipper's Whop USDT position.
- During periods of video plateau where views trickle in slowly, the channel simply stays open longer and settles a smaller amount at the next interval. No spam transactions, no dust, no wasted gas.
- During viral surges, the channel accumulates a larger balance and settles accordingly.

The interval cadence is tuned to balance user experience (clippers see their WLP balance update frequently enough to feel like continuous payment) against on-chain efficiency (batching enough value per transaction to justify the settlement).

22.1.1 Gasless UX and Cost Sustainability

This batching is what makes the gasless user experience economically sustainable for Content Rewards. Every on-chain transaction on Aptos requires gas, and Whop or Content Rewards sponsors all gas costs so that clippers and creators never pay fees; the entire experience is gasless from the user's perspective, funded by gas accounts that Content Rewards maintains and tops up. Without Shelby's micropayment channels, every individual view accrual would be a separate on-chain settlement, and the gas sponsorship cost would scale linearly with view volume. The micropayment channel architecture collapses per-view events into batched settlements, reducing Content Rewards' gas sponsorship burden roughly tenfold. A campaign that would have required 5,000 sponsored transactions under a per-view model might settle in 500 batched transactions over the same period: same USDT delivered to the same clippers, but at a tenth of the gas cost to Content Rewards. As campaign volume grows, the gas account scales with the number of active clippers and settlement intervals, not with the number of views, which means Content Rewards can maintain gasless UX profitably even as the platform scales to millions of daily views across thousands of campaigns.

22.2 Withdrawability

The clipper's full WLP balance (both their principal earnings from clipping and any yield accrued on top) is fully withdrawable at any time with no lock-up, no delay, and no penalty. There is one exception: if a creator opens a dispute on a campaign and a bounty hunter identifies the clipper's submissions as potentially fraudulent, the dispute contract places a lien on the portion of the clipper's WLP balance corresponding to the alleged overpayment. While the lien is active, that portion is frozen (non-withdrawable, non-transferable) until the dispute resolves through the ADR voting process. If voters confirm the fraud, the lien balance is clawed back into the CRIF's Lien Escrow and disbursed to the creator. If voters reject the case and the clipper is cleared, the lien is released and the full balance becomes withdrawable again once the resolution is finalized. Outside of an active dispute lien, nothing restricts a clipper's ability to withdraw everything they've earned at any time.

22.3 Withdrawals

When a clipper wants to access their earnings, they initiate a withdrawal from the WLP. The withdrawal uses Aptos script composer to execute a single atomic transaction that redeems the clipper's Whop USDT from the vault and places the corresponding USDT balance directly into their Aptos wallet. From there, the clipper routes the funds to their destination: transfer to a centralized exchange for fiat off-ramping to a bank account, or direct usage elsewhere on Whop.

Clippers don't need to understand DeFi, vaults, or yield mechanics to start earning; their USDT simply accumulates in the WLP by default like a traditional savings account would, but at significantly higher rates, earning yield in the background. When they're ready to withdraw, the path to a bank account is a single action. The WLP becomes the default financial home for clipper savings after earning, not because of lock-ups or friction, but because leaving money in is more profitable than pulling it out immediately.

22.4 Yield, Voting, and Risk Structure

For Whop users the WLP is a savings account: the place where earnings land and grow automatically at rates that outperform anything a traditional bank offers on idle cash. A clipper earning \$3K/month sees their earnings settle into the WLP automatically, earns risk-adjusted yield from insurance premiums, and participates in ADR votes on disputes in their domain of expertise. They understand clipping, so they can evaluate bounty hunter evidence intelligently. The clipper doesn't think of themselves as a DeFi participant. They think of themselves as someone with funds in a savings account that pays well.

For external depositors: Anyone outside the Whop ecosystem who holds USDT and wants yield can deposit into the WLP, receive Whop USDT tokens representing their share, and earn the same insurance premium yield that clippers earn, without ever clipping a video or submitting to a campaign. External depositors also participate

in ADR voting, contributing to dispute resolution while earning returns on parked capital. The WLP doesn't distinguish between clipper-earned USDT and externally deposited USDT for yield purposes: both contribute to the liquidity pool, both earn the same yield, and both carry the same Tier 2 ADL risk profile. External capital deepens the WLP's liquidity, which strengthens the CRIF overflow, which makes the insurance infrastructure more robust, which makes Content Rewards safer for creators and clippers. That safety is what generates the campaign volume that funds the yield external depositors came for.

Voting: Carries no per-vote financial reward and no penalty for voting on the losing side. Depositors vote because they want the system to work honestly. There is no tiered APY based on voting accuracy, no bonus for picking the winning side, and no punishment for abstaining on a particular case.

Voting recusal: WLP depositors are automatically excluded from voting on disputes involving campaigns they submitted to as clippers. The recusal extends beyond the clipper's own wallet to include connected addresses identified through on-chain wallet graph analysis (wallets linked by direct transfer history, shared funding sources, or cluster patterns consistent with common ownership or close association, the same type of graph analysis used by tools like Bubblemaps). A clipper who submitted to the disputed campaign and whose wallet is connected to three other WLP depositor wallets through regular transfers would see all four wallets recused from voting on that dispute, not just their own. This prevents clippers from coordinating with friends or community members to shift votes, a particularly acute risk on smaller disputes where total voting activity is low, and results can be easily manipulated.

CRIF health signal: When the CRIF is healthy and at target, the WLP earns consistent premium income. When the CRIF is below target and refilling, the premium flow pauses. WLP yield is a direct signal of system health. Depositors can monitor CRIF health on-chain in real time.

Deposit reputation: Sustained WLP participation contributes positively to a clipper's Content Rewards Trust Score. A clipper who consistently maintains and grows their WLP deposit over time accumulates a reputation signal that reflects long-term commitment to the platform's health. The score contribution compounds gradually: consistent deposits over months carry more weight than a single large deposit, because sustained behavior is harder to fake than a one-time capital commitment.

23 Auto-Deleveraging (ADL) and the Loss Waterfall

When a dispute resolves through ADR, the system follows a strict loss absorption waterfall that prioritizes accountability, then insurance, then socialized risk, in that order. Because WLP voters deliver per-clipper verdicts, the confirmed claim amount may be less than the creator's original filing; only the portion attributed to clippers voted guilty constitutes the confirmed claim that enters the waterfall.

23.1 The Backstop Waterfall (Four Tiers)

23.1.1 Tier 1: Lien Liquidation (Principal and Yield Exposed)

- For each clipper voted guilty, the smart contract converts their lien into a seizure.
- The guilty clippers' WLP positions (yield and principal) are clawed back to cover their attributed share of the confirmed claim, effectively liquidating their WLP for dispute settlement.
- The seized funds are transferred into the CRIF's Lien Escrow for disbursement to the creator.
- If a guilty clipper has already withdrawn some or all of their earnings before the lien was placed, the system seizes whatever remains and records the unrecoverable remainder as a DebtRecord on-chain against that wallet, blocking future Content Rewards earnings until the debt is settled.

23.1.2 Tier 2: Insurance Fund Absorption (Insurance Before Socializing Losses)

- If lien liquidations do not fully cover their attributed share of the dispute, the Insurance Reserve draws down to fill the gap.

- The Insurance Reserve only covers lien liquidations from clippers who've withdrawn earnings from WLP, never innocent clippers who are not guilty.

23.1.3 Tier 3: Whop Liquidity Provider ADL (Yield-Only, Principal Protected)

- Triggers only if both lien liquidations from clippers seizures and insurance reserve together cannot cover the confirmed portion.
- The system ranks depositors using a multi-factor auto-deleverage score:
- The shortfall is distributed through precision-calibrated haircuts: the system takes only the exact pro-rata amount each depositor owes, never zeroing any individual depositor's yield entirely unless their proportional share exceeds their total accumulated yield.
- No WLP depositor's principal is ever touched, unless proven guilty.
- Epoch-batched dispute resolution: Disputes that resolve within the same epoch settlement window are netted into a single waterfall pass, preventing sequential processing from compounding honest depositor haircuts.
- Transparency: Every depositor can see their ADL ranking on-chain in real time. Self-correcting voluntary market participants who stay justify risk-adjusted returns and are willing to backstop the insurance fund.

23.1.4 Tier 4: Future Campaign Fee Credits (Final Backstop)

- If total cumulative value across the entire WLP cannot cover an extreme shortfall after offender ADL, Insurance Reserve depletion, and standard ADL, the residual is covered by Future Campaign Fee Credits, allowing for fee-less campaigns that make up for the rest of the creator loss.



Figure 22: ADL Loss Waterfall: four backstop tiers with explicit trigger conditions and priority ordering.

● ADL SCORE

$$ADL = Y_{cum} \times \frac{Y_{cum}}{P_{deposited}} \times W_{time}$$

Y_{cum} = cumulative yield · $P_{deposited}$ = principal deposited · W_{time} = time weight

Figure 23: ADL Score formula: cumulative yield squared over principal, weighted by time—prioritizes levying from depositors who earned the most.

24 Future Campaign Fee Credits (FCFC)

FCFC is the partial or final backstop for scenarios where both the CRIF and the WLP's total usable value have been fully exhausted. This should be exceptionally rare; requiring all the gains from guilty offenders plus total claims larger than the Insurance Fund plus all honest depositors' additional yield (without principal) within an epoch.

24.1 How It Works

- FCFC is scoped to the 7% Content Rewards fee, not the full campaign budget. Content Rewards doesn't own the distribution (clippers earn the campaign budget for generating views). Only the 7% fee the creator would normally pay on future campaigns is credited, effectively making new campaigns free to restart or launch new campaigns without any burden of additional fees.
- FCFC is designed to kickstart campaign usage during the exact moment the platform needs it most. The credit that compensated the creator for a bad experience is the same thing that refills the insurance fund. The system can regain trust while preventing cold start problems.

24.2 Decay Mechanics

- The creator receives FCFC at face value the moment the dispute is resolved.
- The FCFC decays for the equivalent period of time as the disputed campaign ran before being paused.
- If the campaign ran 30 days before the creator filed the dispute, the FCFC has a 24-hour grace period, then decays linearly over 30 days to zero.
- The 24-hour grace period gives the creator time to adjust their campaign before relaunching, to limit disputes by reducing eligibility criteria, increasing the floor CPM, or raising eligibility score.

Creator reputation from FCFC usage: A creator who receives FCFC and fully uses it by launching a new campaign earns a boost to their creator reputation score. A creator who lets their FCFC decay to zero without using it gets no reputation boost; the system only rewards creators who actively reinvest.

Protocol accounting: FCFC is debt on the protocol's books. Content Rewards tracks total outstanding FCFC as a liability and reports it transparently. If outstanding FCFC grows faster than new campaigns retire it, that's a signal to tighten Phase 2 subjective requirement access or adjust the confidence floor upward to reduce dispute volume.

25 Creator Dispute Payout Priority (CDPP)

The waterfall above describes who funds dispute payouts. When multiple creators have confirmed claims resolving in the same epoch and the total recovery pool cannot cover all of them in full, the system must also determine which creators get paid first. The disbursement follows a three-step priority:

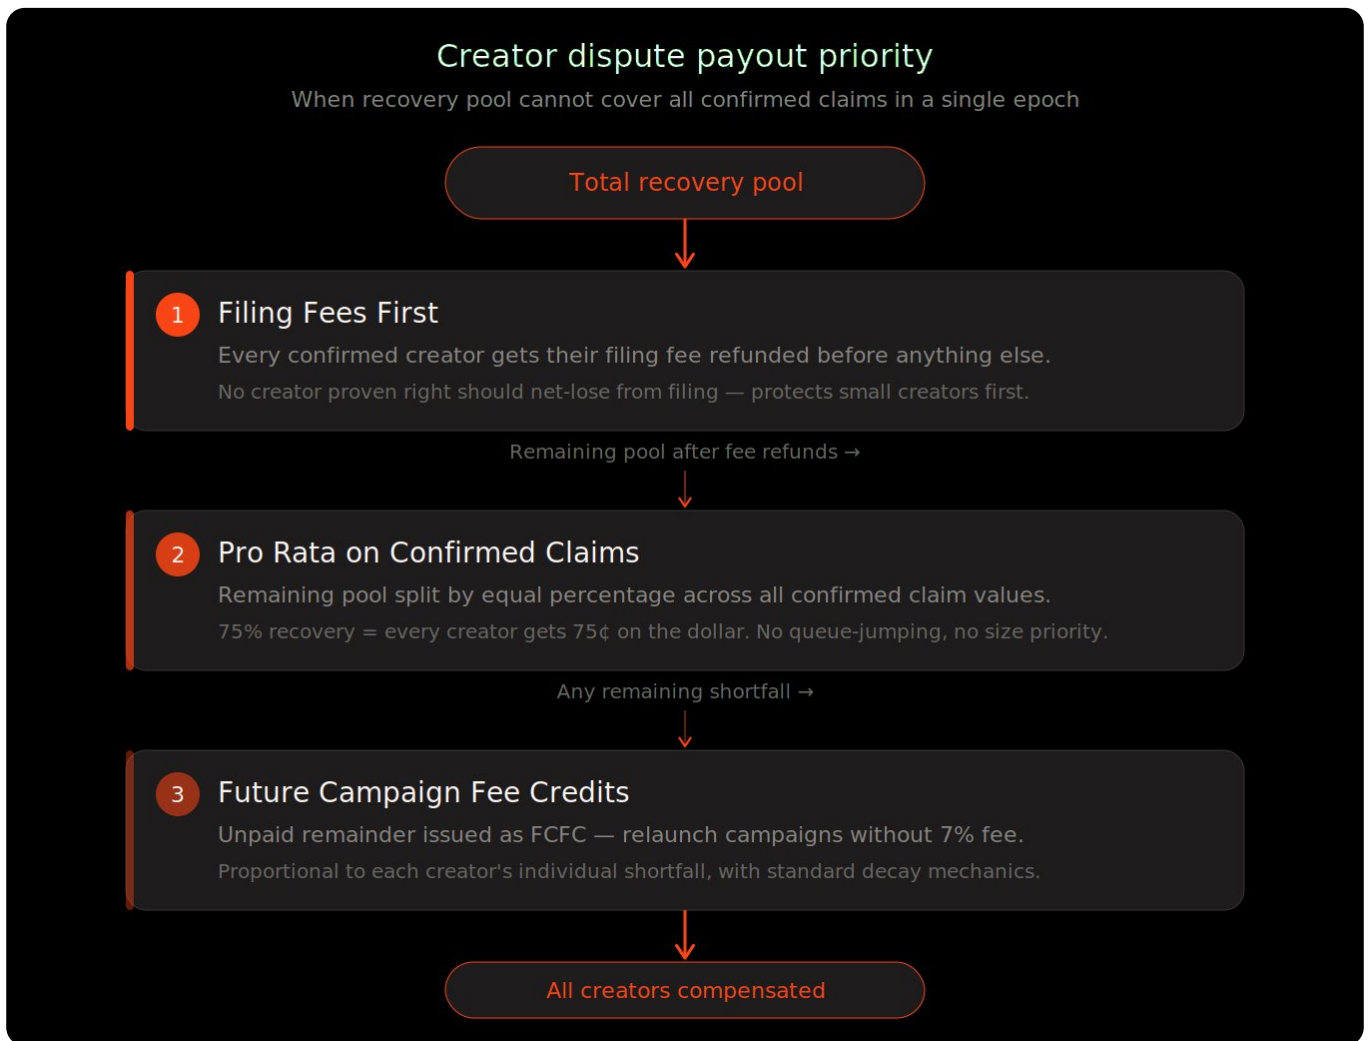


Figure 24: Creator Dispute Payout Priority: three-tier distribution from loss waterfall proceeds.

25.1 The Payout Method (Three Tiers)

25.1.1 Tier 1: Filing Fees First (Entry Costs Refunded)

Every confirmed creator gets their filing fee back before anything else. No creator who was proven right should net-lose from filing; this compensation disproportionately affects small creators first.

25.1.2 Tier 2: Pro Rata on Confirmed Claims (All Creators Reimbursed Equally)

Pro-rata on confirmed claims. Remaining pool split by equal percentage across all confirmed claim values. A 75% recovery rate means every creator gets 75 cents on the dollar regardless of claim size. No queue-jumping, no size-based priority, hardest structure to game.

25.1.3 Tier 3: Future Campaign Fee Credits (Relaunch Campaigns Without Fee)

FCFC covers the deficit. Whatever remains unpaid is issued as Future Campaign Fee Credits proportional to each creator’s individual shortfall, with the same decay mechanics defined above.

26 Roadmap to Phase 2

Phase 2 unlocks slightly more subjective campaign eligibility options that expand the total addressable market to attract quality creators and businesses with larger budgets, but are harder to verify definitively (positive sentiment, strict brand requirements, etc.). Phase 2 activates only when three conditions are met simultaneously:

CRIF threshold: The Content Rewards Insurance Fund (CRIF) exceeds a minimum threshold, with sufficient capital to cover future dispute reimbursements based on projected dispute rates and value.

WLP liquidity: The WLP has attracted enough capital from Whop clippers and external depositors providing market-funded backstop beyond the protocol's own reserve.

ADR track record: The ADR bounty hunter system is operational with active bounty hunters claiming and resolving cases, a working voter pool, and a demonstrated track record of resolved cases validating the mechanism works without internal dispute resolution.

Until all three conditions are met, campaign creation only allows objective requirements that are easily verifiable. Campaign creators are shown: "Subjective requirements are not yet available, please be more specific about eligibility. Content Rewards currently prioritizes instantly verifiable eligibility requirements only to minimize campaign creator disputes with clippers to ensure both parties are in perfect agreement."

If the system never reaches Phase 2: for example, if enough capital flight from WLP before the offender is proven guilty after vote and lien, the insurance fund doesn't fill enough, or the WLP doesn't maintain the minimum long-term depositors, then Content Rewards stays on objective-only eligibility until further notice.

This is not a failure state. A system that guarantees 100% approval on verifiable criteria with streaming USDT payouts, reputation-priced CPM curves, and anti-bot floor calculations is already an improvement over the current antiquated system. Subjective requirements are an expansion plan, not a prerequisite.

27 Summary

PropWhop, an enhanced version of Content Rewards on Whop, replaces every manual bottleneck in the current system (subjective review, delayed payouts, unverifiable metrics) with a market-driven design with programmatic verification, real-time USDT streaming, and an economic structure where fraud is unviable.

Proof-of-results to bypass incumbent gatekeepers through verified results: Clippers convert creative energy into USDT with no bank, no payment processor, and no identity check required. The offramp friction (ID verification at exchanges, bank linking, payout delays) creates a natural retention sink where the path of least resistance is keeping USDT inside Whop, generating yield in the WLP, spending on Whop products and services, or reinvesting into their own community, all P2P and wallet-to-wallet with no KYC.

Campaign creation fee flow: A 20% surcharge on top of the creator's stated clipper budget funds three infrastructure layers: 10% to Whop for operations and USDT onramping, 7% to Content Rewards for management, platform listing, and the CRIF insurance fund (with a dynamic utilization curve that self-balances between operating margin and insurance capacity), and 3% to Aptos for Shelby verification, on-chain settlement, ADR smart contracts, and a self-funding perpetual campaign that markets the platform through its own product.

Content origin verification: Uses dual-channel audio and visual fingerprinting against source material stored as cryptographically committed Blobs on Shelby, with brand-configurable sliders that weight each channel according to what matters for their content type. The 80% confidence floor ensures only clips genuinely derived from source material qualify, with zero human review.

Fraud prevention: Operates through economic incentive design rather than detection-first enforcement. The CPM floor is calibrated below the cost of botting, the Trust Score compounds over a clipper's full history, and account reselling is self-defeating because the environmental signals that accompany an ownership transfer are the same signals that degrade the score.

Real-time verification and payouts: Shelby and Aptos eliminate the review queue entirely. Clippers submit, the system verifies programmatically, and USDT streaming begins immediately. Budget visibility is real-time, creators are notified before depletion, and no clipper is ever denied payment because of queue position.

Trust scores: A clipper Trust Score (1–10) prices access through CPM rates rather than gatekeeping. A creator Trust Score encodes contract reliability, partner quality, and commitment trajectory. A Campaign Trust Score provides real-time health monitoring incorporating creator behavior, CRIF depth, solvency ratio, and clipper quality.

Campaign lifecycle: Optimization actions flow through instant (clipper-favorable) or delayed (clipper-adverse) channels. Sunset campaigns wind down gracefully over 7 days with sigmoid-curved haircuts. End campaigns provide immediate 50/50 splits with steep Trust Score penalties. Dispute campaigns halt immediately with filing fees calibrated to campaign budget.

Automatic Dispute Resolution: Handled through the ADR flow, where creators file disputes, bounty hunters investigate cases and identify fraud, and compile evidence for market-driven voter review. The CRIF operates as two sub-ledgers: an Insurance Reserve funded through a dynamic utilization curve that covers guilty clippers who can't pay, and a Lien Escrow that holds seized funds pending disbursement to creators. The WLP serves as the final backstop through yield-only auto-deleveraging with precision-calibrated haircuts that never touch honest depositors' principal. If financial solvency is in question, future campaign fee credits help kickstart reuse.

Phase 1 launches with objective-only requirements, producing a 100% approval rate for qualifying eligible submissions that unlocks immediate USDT streaming where clippers start earning per view in real time.

Phase 2 unlocks subjective requirements only after the CRIF, WLP, and ADR system all reach maturity simultaneously. If Phase 2 never arrives, the system remains on objective-only requirements until reached.

Powered by Shelby and Aptos, PropWhop is a trojan horse to onboard new and existing clippers on Content Rewards into USDT on Whop with payout guarantees in a familiar product, positioning Whop to follow Shopify's trajectory, with Content Rewards as the entry point and growth driver into Whop Finance.



Content Rewards on Whop

Powered by Aptos and Shelby

Programmatic Verification · Real-Time Settlement · Economic Incentive Design
