SBET V2 · TECHNICAL DUE DILIGENCE

A decentralized sports-betting + prediction-market protocol with confidence-adaptive finalization.

90-second typical settlement · VRF-selected grader panels · bond-escalated disputes resolving in <96 hours. Targeting $10M+ TVL on Base/Arbitrum.

TARGET TVL	GRADER STAKE	FAST SETTLEMENT
$10M+	**$225k**	**90s**
Go-live target, Base + Arbitrum	9 × $25k quorum, matches payout cap	Unanimous + 95%+ confidence tier

The pitch

SBET V2 is a non-custodial sportsbook and prediction-market protocol settling most matches in **90 seconds** while giving disputed outcomes hours of human review. It is the first protocol to tie oracle collateral directly to per-match exposure: a colluding grader quorum stakes ≥ what it can steal, making attacks economically irrational. Ground-up rewrite of V1; four ship-blocking bugs fixed pre-audit; Q2 2026 audit scheduled.

Market opportunity

Decentralized betting and prediction markets captured **>$2B of on-chain volume** in 2025 across Polymarket, Overtime, Azuro, and Thales. Incumbents optimize one dimension at a time: Polymarket + UMA's 2-day dispute window; Overtime's AMM sportsbook without native disputes; Azuro's LP sportsbook with centralized feeds. **None settle in seconds on uncontested outcomes while maintaining decentralized dispute resolution.** SBET V2 targets the intersection: sportsbook + prediction markets, one protocol, 90s-latency on 95%+ outcomes, Kleros-style escalated disputes on the rest.

Differentiators vs incumbents

Dimension	SBET V2	Polymarket / Overtime / Azuro
Finalization latency (uncontested)	**90 seconds** adaptive	Fixed 2h+ windows or post-match-end delays
Oracle trust model	13-grader VRF panel, 9-of-13 quorum, $25k each	UMA optimistic / Chainlink feeds / centralized
Dispute resolution time	<24h (default) to 96h (full escalation)	UMA DVM vote (≥2d) or no native dispute
Per-match exposure cap	$225k (= grader stake)	Pool / LP-bounded, not oracle-bounded
Emergency response	3-tier: match / category / global	Admin pause only

Technical moat

The V2 design is anchored on **13 locked architectural decisions**. Every contract file carries `[Decision #N]` and `[Bug Fix #N]` markers tying source back to the decision log.

- **Confidence-adaptive finalization.** Challenge window scales 90s→4h based on grader unanimity, AI confidence, and source agreement. Uncontested outcomes settle 40–60× faster than competitors.
- **Collateral-bounded exposure.** Per-match cap = grader stake × quorum ($25k × 9 = $225k). A colluding quorum's max gain is ≤ its max loss. No incumbent ties exposure to oracle collateral.
- **6-tier AI grader escalation.** Pipeline climbs cheap→mid→expensive LLMs, multi-source verification, re-grade, human abstention. Operators set per-match spend caps.
- **Bond-escalated disputes.** Initial 1% TVL clamped [$5k, $50k], 2× per round, max 3 rounds → 5-of-9 arbiter. Griefers pay exponentially per round of delay.
- **Three-tier pause + hybrid accounting.** Match (bond, 24h) / Category (3-of-5, 72h) / Global (5-of-9, 14d + 72h timelock). Signed-position + parimutuel in one protocol.
- **Gasless auto-claim (ERC-2771).** 2% SBET / 3.5% other. Trusted forwarder is immutable — rotation requires redeploy, eliminating replay vectors.

Audit-ready posture. OpenZeppelin v5.x (ReentrancyGuard, SafeERC20, AccessControl, Pausable) + strict CEI discipline on every external function. No custom assembly outside loop counters. ~3,000 LOC for auditors: 7 contracts + 2 libraries.

Traction indicators

Milestone	Status	Notes
V1 live on Sepolia	Deployed	~25 kLOC diamond; 13 contract screens in dapp
V2 contracts	Pre-audit	7 contracts, ~130 kB, [Decision #N] markers
Indexer + AI layer	Running on VPS	Postgres + pgvector + Redis; SSE feed
Grader reference impl	Python prototype	6-tier escalation; Go release planned
Q2 2026 audit	Scheduled	Scope: 7 V2 contracts + 2 libraries
V2 mainnet	Target Q3 2026	Base + Arbitrum; V1 settles independently

Risk factors (candid)

These are the risks we actively track. They are not marketing disclaimers — each is a design trade-off we've made explicit, with mitigations in source.

Grader key compromise. Compromised grader can submit wrong grades until arbiter override or dispute slashing. **Mitigation:** per-match VRF rotation bounds single-key damage; economic invariant holds. **Open work:** per-grader key rotation not yet wired (flagged in source).

VRF sequencer influence at L2. Base/Arbitrum run centralized sequencers. **Mitigation:** Chainlink VRF v2.5 seeds are unbiased — sequencer cannot change seed value, only inclusion timing. Blockhash fallback (24h) is miner-influenceable within 256 blocks — accepted as liveness, not security.

Guardian quorum capture. 5-of-9 global guardian compromise can pause the protocol for 14 days. **Mitigation:** 72h unpause timelock prevents mask-attack cycles; pause does not release custody. Guardian diversity is a governance process concern.

Regulatory exposure. Jurisdiction-dependent. V2 is non-custodial; protocol holds no fiat, does not onboard users. **Mitigation:** geographic exclusion at frontend / relay layer without contract changes; treasury governance enforces token allowlists.

Oracle economic ceiling. $225k per-match cap (9 × $25k) is the hard ceiling for individual-match TVL. **Mitigation:** governance can raise stake floor; cap scales proportionally. $10M+ TVL aggregates across matches, not concentrated on single events.

Ecosystem. Four repos: Solidity (V1 Sepolia + V2 pre-audit), React/wagmi dapp (13 contract screens, 5 languages), Python indexer (pgvector + Redis + AI on VPS), docs. V1 dapp already ships Polymarket import, community pools, Kelly sizing, AMM/orderbook routing, maker copilot — V2 inherits all frontend; redesign only touches finalization + dispute.

Next steps for due-diligence readers. Pair this brief with the *V2 Technical Whitepaper* (full architecture, security invariants, parameter appendix) and the *V2 Security Summary* (threat model, worked attack scenarios). Code is available at `github.com/sbettoken/contracts` (path `v2/`).