SBET V2 / SECURITY SUMMARY

Security Summary

Trust boundaries, threat model, defense layers, worked attack scenarios, and known risks for the SBET V2 protocol. Condensed brief for auditors and security-conscious users; pair with the Technical Whitepaper for full detail.

1. Trust Boundaries

What is trusted vs trustless at each layer. Any component marked **trusted** is a security assumption that auditors should challenge; any component marked **trustless** enforces its guarantees via on-chain code or economic incentives.

Layer	Component	Trust	Enforced by
Chain	EVM correctness, block finality	Trusted	Base / Arbitrum sequencer + L1 settlement
Randomness	Chainlink VRF v2.5	Trusted	Chainlink cryptographic proof; fallback bounded
Custody	SBETTreasuryV2	Trustless	CORE_ROLE gate; immutable forwarder
State machine	SBETCoreV2	Trustless	Explicit guards, monotonic state, no custom storage writes
Grader oracle	GraderRegistryV2 + panel	Economic-trustless	$25k stake + VRF draw + slashing
Dispute	DisputeManager + arbiter	Partly trusted	Bond escalation + 5-of-9 Safe for final ruling
Emergency	GuardianCouncil	Trusted (quorum)	3-of-5 category / 5-of-9 global + 72h timelock
Auto-claim	ERC-2771 forwarder	Trusted (immutable)	Set at Treasury constructor; redeploy only

Core holds no user funds. If Core is compromised, user stakes remain protected by the role gate on Treasury. The trusted forwarder is immutable by design — rotation requires redeploy, eliminating signer-rotation replay vectors.

2. Threat Model

2.1 Attacker classes

Class	Capability	Goal
Compromised grader	Controls 1..N grader keys; can submit wrong grades on matches they're drawn onto	Finalize at wrong outcome + capture payouts
Colluding majority	9+ grader keys compromised simultaneously	Force Fast-tier false finalization
MEV / frontrunning bot	Observes mempool; can reorder within block	Extract dispute bounties; sandwich claim txs
Dispute griefer	Can fund up to 8× initial bond ($400k)	Delay finalization; extract bond capital via attrition
Match-lock spammer	Can fund many $1k SBET bonds	Censor pause many matches simultaneously
Guardian partial-compromise	1–4 global guardians, 1–2 category guardians	Cannot reach quorum; observable on-chain
VRF fallback miner	Controls blockhash within 256 blocks	Pre-compute favorable panel after 24h VRF delay

2.2 Assumptions (explicit)

- **EVM correctness** — underlying chain executes bytecode correctly and finalizes blocks.
- **VRF soundness** — Chainlink VRF v2.5 produces unbiased randomness.
- **Guardian distribution** — 5-of-9 guardians and 5-of-9 arbiter multisigs are not *simultaneously* compromised.
- **Grader economic rationality** — graders act to maximize expected return on their $25k stake.
- **Bondholder liquidity** — challengers and defenders can source bond capital on demand.

3. Defense Layers

V2 composes six economic and mechanical defense layers. Attacks must defeat multiple layers to succeed; defeating one layer reveals the attack early enough for the next layer to respond.

3.1 Stake + slashing

Grader stake (minimum $25k SBET, held in GraderRegistryV2) is the collateral for honest behavior. Slashed on arbiter ruling; funds flow 50/20/20/10 to challenger / treasury / stakers / bounty. **No burn.** Invariant: `perMatchPayoutCap() ≤ stakeFloor × quorumM` — colluding quorum's max loss ≥ max gain.

3.2 VRF panel rotation

13 graders drawn per match via Chainlink VRF v2.5, quorum 9. Attackers cannot pre-select which matches they grade. Panel membership committed immutably via `panelRoot` (Merkle) at `registerMatch`. Fallback to `blockhash(matchCreationBlock)` after 24h is documented as a liveness — not security — guarantee.

3.3 Challenge window tiers

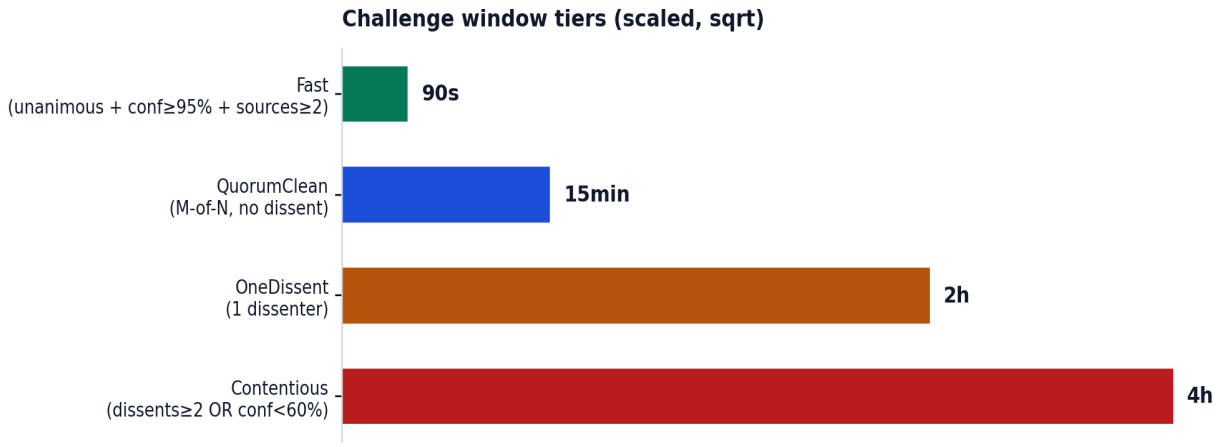

Window length scales with consensus quality: 90s when unanimous + 95%+ confidence + 2-of-3 independent sources agree; 4h when any dissent or low confidence. A rational attacker must open a dispute before `unlockAt` regardless of tier — shortening the window only reduces delay for uncontested matches.

3.4 Dispute bond escalation

Initial bond = 1% match TVL, clamped [$5k, $50k]. Each of 3 rounds doubles stake. Max cumulative 15× initial. Missing a 24h response → `claimByDefault` forfeits contributions to winner. After round 3, 5-of-9 arbiter rules.

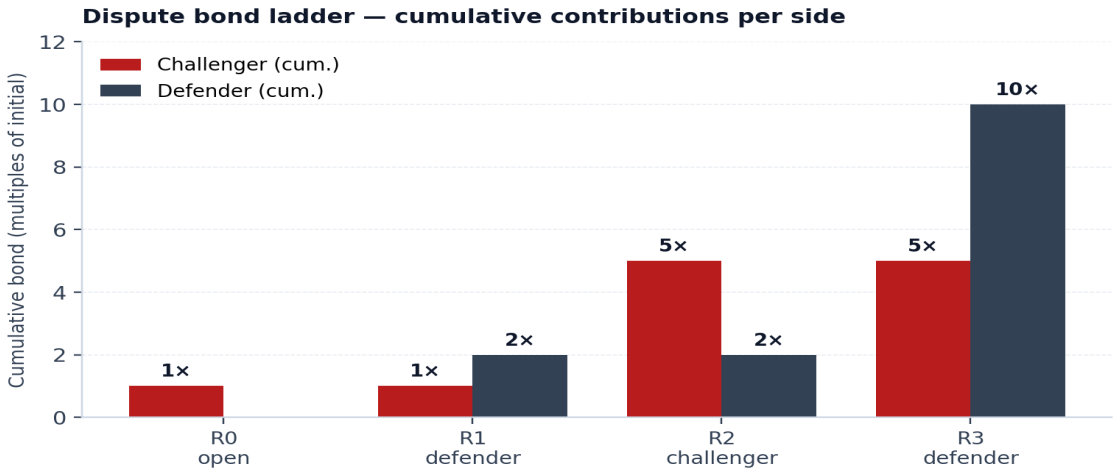

3.5 Guardian multisig + three-tier pause

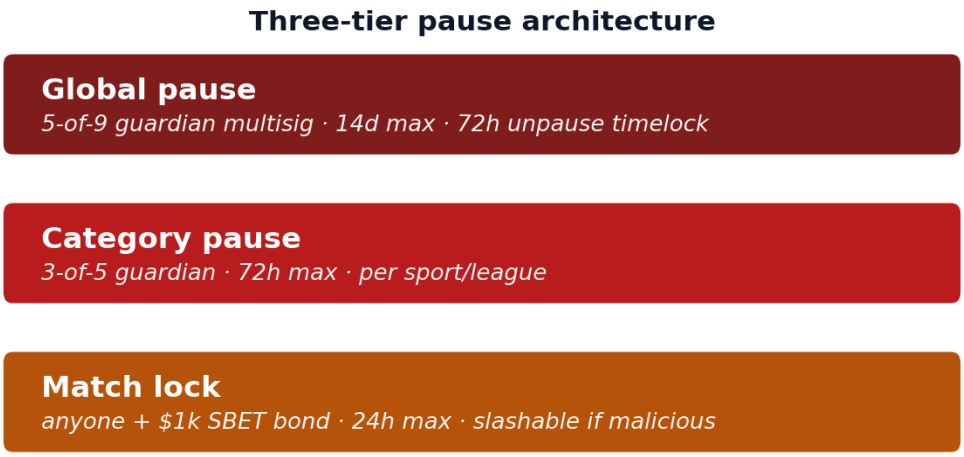

Tiers compose. Active tier = highest-scoped active pause.

Scope-appropriate responses: permissionless match-lock (24h, $1k SBET bond) for individual incidents; 3-of-5 category pause (72h) for sport/league-wide issues; 5-of-9 global pause (14d) for protocol-wide emergencies, with 72h unpause timelock to prevent mask-attack cycles.

4. Worked Attack Scenarios

S1 — 9-of-13 grader collusion on $225k match

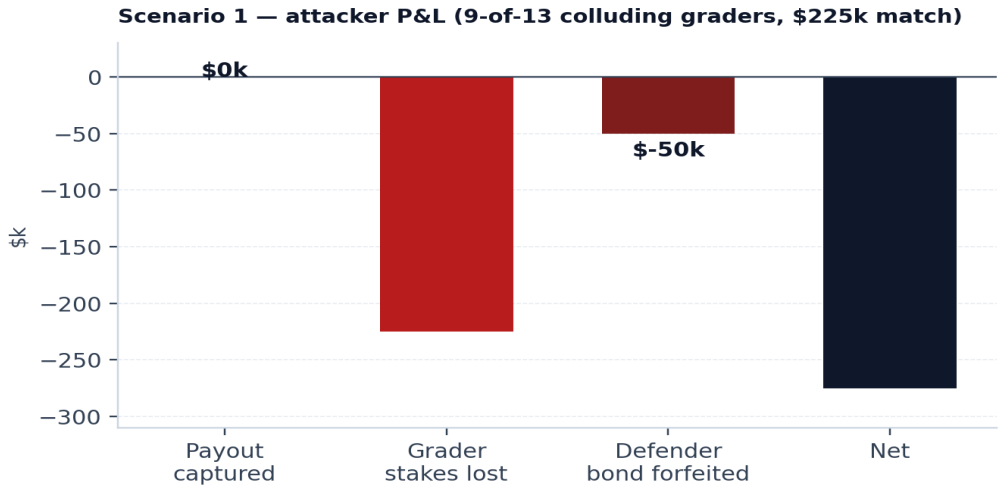

Attacker controls 9 grader keys ($225k total stake). Targets a $225k TVL match. 9 signers post unanimous wrong outcome, confidence 10,000, sources 3 → Fast-tier 90s window. Honest $5k bond dispute opens within 90s. Bond escalates to $75k at round 3. Arbiter rules for challenger.

Attacker P&L: payout captured = $0 (dispute prevented finalization); grader stakes slashed = -$225k; defender bonds forfeited to challenger = -$50k. **Net: -$275k.** Challenger earns ~$112.5k (50% of slash) + bond refund. **Invariant holds:** max attacker loss ($275k) > max gain (capped $225k).

S2 — 90-second censorship race

Attacker wants Fast-tier window to elapse before any dispute is opened. A $5k hot-wallet monitoring bot can call `openDispute` in <10s. Defeating this requires network-wide DoS of every monitoring bot for 90 seconds. Even on successful censorship, 5-of-9 arbiter can `overrideOutcome` on {Open, Proposed, Disputed}.
Post-finalization recovery requires new governance action (not yet wired — flagged as open question).

S3 — Dispute griefing to delay finalization

Griefer targets $10M TVL match. Initial bond = $50k (ceiling). Griefer opens ($50k). Defender posts $100k. Griefer must post $200k round 2 — defaults. `claimByDefault` → griefer's $50k forfeited; resolution propagates at defender's outcome. **Amortized cost: $50k per 24h delay = 0.5% of TVL.** Positive protocol P&L on each griefing attempt.

S4 — Match-lock spam

Attacker locks 100 matches at $1k SBET bond each = $100k capital. Governance slashes malicious locks via `releaseMatchLock(matchId, true)`; bonds flow to Treasury bounty pool. **Attacker loses $100k; protocol gains $100k.** Only attacker upside: 24h delay per match if governance does not slash. Bounded damage, positive protocol revenue on slashing.

S5 — VRF fallback miner collusion

VRF does not fulfill within 24h. Miner controlling `blockhash(matchCreationBlock)` pre-computes a panel they've compromised. **Mitigations:** (1) matches with `VRFSeedFulfilled.fromFallback=true` are on-chain visible — risk engines can flag; (2) 4h contentious challenge window gives dispute openers more runway; (3) scenario S1's economic bound still holds — attacker loses $225k if dispute raised. **Residual risk:** acceptable as a liveness guarantee.

5. Known Risks & Open Issues

Flagged in V2 source for follow-up. Items below are targets for the Q2 2026 audit scope and subsequent governance action.

5.1 Open questions flagged in source

- **On-chain arbiter vote aggregation.** ARBITER_ROLE trusts the 5-of-9 Safe multisig to enforce quorum off-chain. Adding an on-chain vote phase would make quorum visible on-chain at the cost of more complex state.
- **Dispute pot distribution.** Currently dispute bonds go 100% to the winner. Future revision may route through the 50/20/20/10 split to keep economic alignment consistent with slash distribution.
- **Grader slashing on arbiter ruling.** DisputeManager.arbiterResolve has a TODO for triggering GraderRegistryV2.slashWithChallenger. Requires DisputeManager to hold SLASHER_ROLE and a per-match signer-list view.
- **VOID fee-rebate schedule.** claimVoidRefund currently refunds principal only; per-PoolType fee schedule not yet locked.
- **Per-grader key rotation.** Not yet wired; requires governance path distinct from stake lifecycle (register / deregister).
- **Post-finalization recovery.** If a match finalizes incorrectly (e.g. due to arbiter error), there is no canonical re-open path today.

5.2 Accepted residual risks

- **VRF fallback blockhash bias** — miner-influenceable within 256 blocks. Accepted as liveness, not security.
- **L2 sequencer censorship** — Base/Arbitrum sequencers are centralized. SBET inherits L2 risk; no SBET-specific defense.
- **Guardian multisig compromise** — 5-of-9 compromise pauses protocol for 14 days. Cannot drain funds. Mitigated by 72h unpause timelock and quorum requirement.
- **Arbiter ruling finality** — no on-chain appeal against arbiter rulings today. Off-chain governance can re-open via override (role-gated).

5.3 Audit scope (Q2 2026)

Seven contracts + two libraries, ~130 kB source, ~3,000 LOC including interfaces. Priority areas:

- SBETCoreV2 state machine — all 8 transition paths, role checks, reentrancy
- SBETTreasuryV2 — slash distribution 50/20/20/10 invariant, ERC-2771 forwarder immutability, auto-claim fee accounting
- GraderRegistryV2 — stake lifecycle, aggregation math, slashWithChallenger role gate
- DisputeManager — bond escalation ladder, claimByDefault paths, arbiter resolution
- ChallengeWindowLib + SignedPositionLib — pure math libraries, overflow / precision
- GuardianCouncil — pause quorum + timelock logic, composition of the three tiers

- Cross-contract calls — CEI compliance, reentrancy across Core → Treasury → StakerRewards chain
- Formal invariants (I1–I8) as fuzz/symbolic test targets

6. Bug Bounty Scope (Outline)

Post-audit, V2 deployment will include a bug-bounty program. Preliminary scope and severity tiers below; final program details will be published at deployment.

Severity	Example	Indicative reward
Critical	Direct custody theft; bypass of per-match cap; arbitrary state mutation bypassing role gate	TBD (uncapped, %-TVL-scaled)
High	Slash invariant violation; incorrect payout beyond rounding; panel spoof	TBD
Medium	DoS on claim path; unintended griefing vectors; event misreporting affecting off-chain indexers	TBD
Low	Missing NatSpec; gas inefficiency; docs / source mismatch	TBD

In scope

- All seven V2 contracts + two libraries (SBETCoreV2, SBETTreasuryV2, GraderRegistryV2, DisputeManager, GuardianCouncil, SBETStakerRewards, ChallengeWindowLib, SignedPositionLib)
- Interfaces in `contracts/v2/interfaces/`
- Protocol-level economic invariants (I1–I8)

Out of scope

- V1 contracts (separate scope under existing V1 audits)
- Dapp frontend / indexer / AI pipeline (separate review tracks)
- Off-chain grader operator infrastructure
- Third-party libraries (OpenZeppelin v5.x, Chainlink VRF — covered by upstream audits)
- L2 sequencer / Chainlink VRF service-level behavior

Responsible disclosure. Submit findings privately to the Security WG before public disclosure. Post-audit contact details and PGP keys will be published alongside the Q3 2026 mainnet deployment. Do not test on mainnet in ways that affect live user positions.