VMark: A Market for Cyber Risk Settled on Proof of Capability

Omar Bohsali
omar@vmark.com
28 August 2026

Abstract. Cyber risk is unique among major risk classes: the capability to cause a loss can be demonstrated without causing it. Credit default swaps pay after default, parametric catastrophe contracts pay after a physical threshold is crossed, and insurance pays after a claim. We propose a market that uses this proof of capability as the settlement event. An operator lists a named system, which we call a realm, thereby authorizing in-scope, non-destructive tests of published types. Participants take long or short positions on whether a valid proof against that realm will be verified before a stated time. The party that can produce the proof takes the short, notifies the operator privately, and is paid when the proof is verified. No separate finder's fee is used. The same instrument pays for a hole in one system and for a zero-day in a component shared by many: a researcher who finds a widely deployed flaw can take the short on every listed realm on which it can be demonstrated, then submit. The resulting price is whether a qualifying proof will be verified, not expected loss. As long as only in-scope, non-destructive proofs settle, a researcher ought to find it more profitable to trade the proof than to cause the corresponding harm.

1. Introduction

Cyber defense is presently organized around prevention, after-the-fact insurance, and vendor-run bounty programs. While these work well enough for many incidents, they do not produce a price for residual exposure on a named system. Bounty payments are set by the vendor. Insurance pays after loss and typically excludes the correlated events that matter most. Purchasers of vulnerabilities on illicit markets do not publish prices. Recent proposals for a collective response to AI-enabled attacks have emphasized tools, funding, and coordination [1]. They do not specify a settlement event that would allow capital to move before the loss.

What is needed is an instrument that can be priced continuously and that pays on evidence that a loss was available, not on the loss itself. In this paper, we propose such an instrument. We rely on a property of digital systems that other major risk classes do not share: capability can be shown without the act. We call this Rule 1. The market we describe has one contract, two locked pools, and no leverage. Settlement is a verified, non-destructive proof. The pools are parimutuel because the researcher may already know the answer.[1] VMark is a centrally operated market, not a cryptocurrency protocol.

2. Proof of Capability

We define a proof of capability as a non-destructive, independently reproducible demonstration that a specified attack would have succeeded on a listed realm. The proof must not itself cause the harm it demonstrates. Ordinary instruments cannot use this kind of evidence. A credit default swap cannot pay because a borrower could have defaulted; it pays because they did. A catastrophe bond cannot pay because a hurricane could have formed. In cyber, the analogue of default can be exhibited in advance. Rule 1 is this fact: cyber is the only major risk class in which the loss can be proven possible without the loss occurring. That holds where a non-destructive marker can stand in for the protected action.

The marker must represent the contracted capability, not a nearby harmless act. Permission to write only a reserved DNS name is not proof of control over the zone. A request that can fire a beacon is not proof of remote code execution. The private disclosure has to show that the same unauthorized path could have done the protected thing.[2]

Earlier designs settled on a different event. Exploit derivatives pay if an exploit occurs [2]. Bug bounties pay an amount chosen by the vendor. Insurance-linked securities pay after insured loss. Equity-factor constructions infer cyber risk from the prices of other assets. None of these settle on a non-destructive proof against a named realm.

The first market should start with markers anyone can check. The cases below are those.

3. Examples

Each case names what is protected, what the researcher publishes, and what does not count as a proof.

3.1. DNS zone

An operator lists a domain, say example.com, and names the records that would actually matter if they changed: the zone apex, www, mail exchangers. A researcher who can write the zone does not touch those records. They write a signed TXT at _vmark.example.com containing the contract, a nonce, and the hash of a private disclosure. Anyone can query it: dig TXT _vmark.example.com.

The lookup shows that the zone can be written. It does not show that the path was unauthorized. The disclosure has to show that the same credentials or exploit could have changed a named protected record, and that they did not. An employee using the official registrar console to write _vmark produces

[1] Extra informed capital must change the price against itself. A market maker who kept selling at the old quote would be the other side of a forced settlement. Identities are in Appendix B.

[2] Appendix A calls this capability equivalence and writes each contract as a versioned 7-tuple rather than a loose label such as "RCE."

the marker and fails the proof. A listing may use a reserved CNAME instead of TXT; the test is the same.

3.2. Signing key

An operator lists a named key: a wallet, a code-signing key, a TLS private key. Unauthorized use of that key would be the loss—funds moved, malware shipped, certificates issued. The researcher does none of those things. They sign a message that exists only for this contract: the realm, the tenor, a nonce, and the hash of the disclosure. Anyone who has the public key can check the signature. Funds stay put. Nothing is issued. Production state does not change.

A signature produced under the operator's dual-control procedure is not a proof. A signature over some other message is not a proof. The nonce binds the claim to this contract, so a signature harvested from elsewhere does not settle. The disclosure has to show that the key was used through an unauthorized path.

3.3. Restricted file

An operator lists write access at a named host, container, or privilege boundary—the kind of access that, used fully, would be ransomware or a webshell. The researcher writes a small signed file at a reserved path, for example /var/lib/vmark/proof.txt, or an object in a bucket that should not have been writable. The file holds the contract, the nonce, and the hash of the disclosure. It is removable. It does not encrypt anything, and it does not leave that path. Anyone who can fetch the path can see that it landed.

Write access to a world-writable scratch directory is not remote code execution. The disclosure has to show that the method crossed the contracted boundary, not a permission that exists only to write the marker.

3.4. Web origin

An operator lists a production website: the pages and responses a customer actually gets. The researcher does not deface it. They place a reserved, inert change at that origin. Three forms are equivalent if the listing names them:

- A hidden element that carries a signed nonce and is not shown to users.
- A 1-by-1 pixel, or another reserved URL, that requests a callback bound to the contract.
- A response header that is not otherwise used, for example X-VMark, whose value is the signed payload.

Anyone can fetch the page or the headers and look. The visible customer experience does not change. A server-side request that can fire a beacon from some other host is not proof of origin write. A compromised analytics tag that can load a pixel is not proof of control over the site unless that tag is the contracted boundary. The disclosure has to show that the method used the same content or deployment authority that protects the origin.

3.5. Canary secret

An operator lists read access within a named secret boundary: a vault path, a production database role, an identity-provider secret. They place a rotating canary in that boundary and nowhere else. The researcher authenticates or calls back using the current canary. They do not dump customer secrets. They do not use any other secret. The callback is checkable. The canary rotates, so a copy taken last month does not settle this tenor.

Presenting a canary that leaked through an authorized channel is not a proof. The disclosure has to show that the canary was obtained through the unauthorized method, and that the same method could have read a protected secret in the same boundary.

3.6. Other surfaces

The same pattern applies wherever a small, attributable state change can stand in for the protected action. A listed contract can require a 0-value call or a signed event rather than a drain. A listed network can require a beacon from an address in that block rather than ransomware. A listed account can require a reserved post rather than a takeover. Those can wait until the marker is as easy to check as a public DNS record. The first market should list the cases above.

4. Listing

A realm is a named digital surface an operator is willing to have priced: a domain, a cloud organization, a product, a network, or a combination of these. The operator must prove control of it. Listing is authorization only for the assets the operator names, and only for the published proof classes. Proofs of those types, against the listed surface, within stated bounds, are in-scope tests. Third-party systems and data remain out of scope unless their operators separately authorize testing. Proofs outside those bounds do not settle and are not authorized by the listing.

A realm cannot be listed by a third party. A government, laboratory, or insurer may put up capital, but it cannot authorize a test on someone else's system.

Operators list because the alternative is the status quo, in which a party who has found a hole has no legal book on which to trade that information at a market price. Listing also gives the operator private notice of a valid proof and an opportunity to remediate before details are public. Systems that remain unlisted continue to be exposed to the residual market that already exists; the listing does not create that market.

Under-resourced operators, such as some hospitals and water utilities, can list with a sponsor on the other side of the book. The sponsor's capital is what makes looking at that realm worth a researcher's time. This is distinct from a grant. The contract pays if a proof is produced.

5. The Contract

For each listed realm, proof class, and tenor T, there is one contract: a valid proof of capability against the realm will be verified before T. There is no second instrument for researchers and no administered finder's fee.

The contract has two fully collateralized pools. Let S be short capital and L be long capital. If a proof is verified, the shorts get their capital back and divide L. Each dollar short earns L/S. If no proof is verified before T, the longs get their

capital back and divide S. Each dollar long earns S/L. Capital, once committed, remains locked until settlement or expiry.

Suppose the long pool contains $1,000,000 and the short pool contains $100,000. If a proof lands, each dollar short earns $10. A researcher who puts up $100,000 earns the $1,000,000 long pool. If no proof lands, each dollar long earns ten cents. The pool ratio $S/(S+L)$ is the price: 9.1 percent in this example, before fees.[3] That number is the market's price of a verified proof before T on this contract. It is not the probability of a breach and not expected loss.

More short capital lowers the return to every short. If another $900,000 joins the short pool, S rises to $1,000,000 and each dollar short earns $1 instead of $10. The informed participant changes the price against themselves. No market maker is required to keep selling at the old price. The researcher need not put up all of the short capital. Information and money can be different parties, as in other markets.

Anyone may take either side, subject to the market's trading rules. A researcher who knows a class of devices, an insurer with claims experience, and an operator who has just rotated a key all enter the same pools.

VMark operates the market directly. It holds the collateral, records the positions, closes the books, checks the proofs, and pays the winners. A claimant also posts a separate bond, so that freezing books is not free.

6. Settlement

The steps are as follows:

1. The researcher prepares a disclosure document D describing the method, computes its hash, and prepares the public marker.
2. The researcher sends VMark one instruction containing the short positions, the hash of D, and the claim bond. Either all of the positions are accepted or none are. VMark closes every named book before any realm is notified.
3. The researcher publishes the signed marker and immediately sends D privately to VMark.
4. The proof is checked. Did the committed method cause the marker? Was it in scope, non-destructive, and reproducible? If so, the contract settles and the short pool is paid. If not, the contract reopens and the researcher's short capital remains at risk through T.
5. Exploitable detail remains with the operator. The market is informed that a proof of a given class was verified, not how to repeat the attack.

A proof that causes harm is invalid. A proof outside the listed scope is invalid. A ransom note is not a proof. Exfiltrated data is not a proof.

This single instruction limits the interval between position and disclosure. It also stops the first notice from giving away the rest of the trade. If one flaw reaches many listed realms, the books close together and the operators are notified together. Some delay remains possible before the instruction reaches VMark, as it does in ordinary vulnerability research. The alternative is not always immediate disclosure. The same fact may otherwise be hidden, sold, exploited, or never found. The market works if prompt disclosure pays best.

[3] Write $\hat{p} = S/(S+L)$. This is the break-even probability for a risk-neutral dollar at that pool state. It is not a forecast of breach or of loss. See Appendix B.

7. Verification

A compromise is, to most observers, invisible. That is why cyber risk is hard to price: the loss can be available for months without anyone outside the operator—and often inside it—having a checkable fact. Rule 1 supplies that fact.

Many proof outcomes are publicly observable. Anyone can query a DNS TXT record. Anyone can request a public page and look for a reserved header, a hidden element, or a 1-by-1 pixel that calls back with a nonce. Anyone can fetch a marker file from a misconfigured bucket. A lookup verifies that the marker landed. It does not, by itself, establish that the path was unauthorized or that the committed method caused the outcome. The average person will not make that lookup. Most people do not read certificate transparency logs either. The market does not require that they do. It requires that the evidence be checkable by anyone who cares to check.

There are two things to check. First, the public marker shows the outcome: the record was written, the file appeared, or the key signed. Second, someone still has to check that the committed method caused it. VMark does that check, with the operator's logs and, where the listing says so, an independent laboratory. A financially interested operator cannot decide the claim alone. The first market should prefer proof types where the outcome is public and the method is easy to confirm.

8. Incentive

Markets of this form are reflexive: the existence of the payout changes the probability of the event [3]. If the event were the harm itself, the market would pay for causing it. A market that paid on the date of a person's death would pay the person who chose the date [4]. That is the objection usually raised against trading on disasters or crimes.

We choose a different event. Settlement requires a non-destructive proof on a listed realm. The party best able to predict the event is the party able to produce the proof, and producing the proof is disclosure, not attack. The cheapest way to be paid should be the legitimate demonstration. If some cheaper illegitimate path exists, the definition of the proof is wrong and should be tightened.

A party who has a valid proof ought to find it more profitable to take the short and submit the proof than to sell the same fact into an illicit market or to carry out the corresponding attack. Exploitation is operationally costly and legally exposed. It remains possible, as it is today. The market does not need to prohibit it in order to compete with it. It needs to pay.

The payout has two scales. Against one listed realm, a valid proof settles that contract. That is the incentive to look at a particular hospital, water plant, or firm. Against a shared component—a library, a VPN concentrator, an identity provider, a cloud control plane—the same flaw can be demonstrated on every listed realm that depends on it. The researcher commits the short positions as one portfolio and submits the proofs together. Disclosure is simultaneous. The

payout is the sum of the long pools. A vendor bounty pays once, at a price the vendor set. A widely deployed zero-day, on this market, is a position in every listed realm on which the flaw can actually be demonstrated.

Bug bounty programs remain useful for vendors who wish to buy patches for software they ship, at prices they set. They do not price residual exposure on a named realm, and they do not let a third party take the other side. The two can coexist.

9. Subsequent Contracts

After settlement of a tenor:

1. The operator holds the proof privately.
2. Shorts have been paid.
3. The operator remediates, or does not.
4. The next contract on the same realm prices the new state.

If the condition has been removed, parties who believe that commit long capital to the next tenor. If the proof still reproduces, parties who can still produce it commit short capital. The new pool ratio is what an operator, insurer, or regulator can observe. No second contract on whether the operator will patch is required, and no committee scores remediation. An operator who wants a lower short ratio has to earn it on the next tenor.

A researcher who was short into the proof may commit long capital if they believe the fix, or remain short if they do not. The same construction prices discovery and, later, whether the fix held.

10. Objections

An insider could write the DNS record. An authorized administrator making an authorized change is not a proof. An employee bypassing a published dual-control rule may be a proof, if that rule is in scope. An employee who takes a position and then creates the hole is manipulating the market and does not settle. The disclosure and the logs distinguish these cases. The public marker alone does not.

Won't this pay people, and agents, to find vulnerabilities? Yes. That is the purpose of the instrument.

Who takes the long side? Operators, insurers, governments, and other sponsors can put long capital behind realms they want examined. A trader who believes no valid proof will arrive also goes long. The short pool is the return for being right. If a hospital cannot post long capital, a sponsor's job is to post it. If no one will, the realm has no funded reward and no meaningful market. That fact is visible.

Is the long pool merely a bounty? From a sponsor's perspective it is a refundable, market-priced bounty. If a proof arrives, the pool pays for discovery. If none arrives, the sponsor recovers its capital and earns the short pool. The amount paid to each short is not administered. It is determined by the ratio of opposing capital. The researcher must stake capital, risks rejection, and competes with every other short. A sponsor funds the question; the market prices both answers.

Does the ratio measure cyber risk? It measures the price of a verified proof under one contract. A low price may mean a hard realm, a thin long pool, little attention, or an expensive proof. It is not expected loss.

Can a false claimant freeze the market? Closing books requires a claim bond. Frivolous or fabricated claims put that bond at risk. If the proof is rejected, the books reopen.

Won't this produce more zero-days? Yes. Better they are found here, under Rule 1, than stockpiled. Franklin wrote that three may keep a secret if two of them are dead [5]. The grey market keeps the secret. This market does not.

11. Conclusion

We have proposed a market for cyber risk that settles on proof of capability rather than on loss. We started from the observation that digital harm can be shown without being done. Operators list realms, thereby authorizing in-scope tests against assets they control. Each contract has two locked pools. A valid proof pays the shorts; expiry without a proof pays the longs. The researcher is paid by the book, not by a fee. A shared vulnerability is a short across every listed realm on which it can be demonstrated. Subsequent contracts price whether the condition remains.

The system can begin with a small number of volunteer realms and with proof types that are easy to verify without touching production, such as DNS control and cloud misconfiguration, using designated verifiers and ninety-day tenors. VMark can hold the collateral, close the books, and run settlement directly. The market depends on Rule 1, on a public marker joined to a private disclosure, and on one book that pays for a proof without paying for the harm.

A. Proof Classes

A proof class is a versioned settlement standard, not a loose label such as "RCE" or "DNS attack." A contract is an instantiation

$$C = (R, P, B, M, E, V, T) \tag{A.1}$$

where R is the listed realm, P the exact protected object, B the authority boundary an unauthorized party must cross, M the marker predicate, E the required private evidence, V the verifier, decision rule, and appeal path, and T the expiry.

Two contracts that share a colloquial name are not the same contract. Write-access to a reserved DNS name is not write-access to the zone. A beacon that can be fired from a server-side request is not remote code execution. The listing has to name P and B before anyone trades.

A valid marker is bound to the contract, signed by the claimant, small, removable, and publicly observable where feasible. Its placement must obey the listing's limits on load, availability, data access, and customer impact. The private evidence must make the method reproducible, or forensically checkable if reproduction is no longer possible.

The proof must also satisfy *capability equivalence*. The marker may be harmless. The authority used to place it must be materially equivalent to the authority needed for the protected action. Write A_M for the authority that placed the marker and A_P for the authority that would perform the protected action. Settlement requires that A_M is sufficient for A_P on the listed object, not merely that M became true.

Permission to write only _vmark is not A_P for the zone. A forged request that can call back with a nonce is not A_P for remote execution. A valid signature is not, by itself, A_P for bypassing a key-use policy. The private bundle E is what joins the marker to the unauthorized path.

The examples in the text correspond to versioned classes DNS_ZONE_WRITE_V1, SIGNING_KEY_USE_V1, RESTRICTED_FILE_WRITE_V1, WEB_ORIGIN_WRITE_V1, and CANARY_SECRET_READ_V1. A later revision is a different class. The first market should list only those five, where the marker is public and the method is easy to confirm.

B. Pool Identities

Fix a contract. Let $S > 0$ be locked short capital and $L > 0$ be locked long capital. There is no leverage and no market maker. Capital remains locked until settlement or expiry.

If a proof is verified, each short dollar is returned and receives its share of the long pool. If none is verified before T, each long dollar is returned and receives its share of the short pool:

$$\pi_S = \frac{L}{S}, \qquad \pi_L = \frac{S}{L}. \tag{B.1}$$

The pool-implied settlement price is

$$\hat{p} = \frac{S}{S+L}. \tag{B.2}$$

This is the break-even probability for a risk-neutral marginal dollar at that pool state. A trader who assigns probability q to verified proof has positive expected value on the short side when $q > \hat{p}$ and on the long side when $q < \hat{p}$. The identity does not say that $\hat{p}$ is the probability of a breach, or expected loss, or residual cyber risk in any other unit. It prices whether a qualifying proof will be submitted and verified under the stated scope, funding, attention, and procedure.[4]

The construction is parimutuel because the party with the best information may also be able to force the event. Extra informed capital must therefore move the price. If the existing short pool is S_0 and a proof holder atomically adds $x > 0$, the holder's share of the long pool is

$$\frac{x}{S_0 + x} L \tag{B.3}$$

and the new implied price is

$$\hat{p}' = \frac{S_0 + x}{S_0 + x + L}. \tag{B.4}$$

The original shorts keep $S_0/(S_0 + x)$ of L. They are diluted. That is the cost of being early without being the party who submits, and the reason a market maker is not asked to keep selling at $\hat{p}$.

The numerical example in the text is $L = \$1{,}000{,}000$, $S_0 = \$100{,}000$. Then $\hat{p} = 9.1\%$ and $\pi_S = 10$. If the proof holder adds $x = \$900{,}000$, (B.3) pays that holder \$900,000 of the long pool, every short dollar earns \$1, and $\hat{p}'$ rises to one half. The holder need not supply all of x. Researchers may work with trading firms or other capital providers. VMark does not prescribe that relationship. Information and money can be different parties.

[4] Fees, if any, should be stated as a share of the winning pool before $\hat{p}$ is quoted. The identities here are before fees.

Across n listed realms on which the same capability can actually be demonstrated, a holder who atomically commits (x_i) receives

$$\sum_{i=1}^{n} \frac{x_i}{S_{i0} + x_i} L_i \tag{B.5}$$

if every named proof is verified, and otherwise keeps the short capital at risk through each tenor. A vendor bounty pays once. This sum pays once per listed realm on which the proof class is satisfied. Realms on which the capability cannot be shown do not contribute, even if they share the same underlying library.

C. Settlement

Let D be the disclosure document and $H = h(D)$ a collision-resistant hash. The claimant prepares D, computes H, prepares the marker, and posts a claim bond. The claimant then sends VMark one instruction: the short positions, H, and the bond. Either every named position is accepted or none is. VMark closes every named book before any realm, any outside verifier, or any other participant is notified.

Only then does the claimant publish the signed marker and send D to VMark. VMark gives the operator, and where the listing requires it the designated laboratory, either a copy of D or the minimum the precommitted procedure needs. VMark checks, in order: that $h(D) = H$; that the marker satisfies M and is bound to this contract; that the method is in scope; that capability equivalence holds; that the committed method caused the marker; and that the safety limits were kept.

If the proof is valid, the shorts are paid and VMark publishes a signed ruling that names the contract, the proof class, and the outcome, without the exploit. If it is invalid, the book reopens and the claimant's short capital remains at risk through T. All or part of the separate claim bond may be kept, under a published schedule, for claims that are fraudulent, frivolous, or materially outside the listed requirements.

The claim bond is what makes freezing books costly without giving anyone an advance look at D. Its amount, and the maximum review period, scale with the number and size of the books named in the instruction. A claim that is not decided within the published period goes to the listing's appeal authority.

The single instruction has two jobs. It limits the interval between the informed position and disclosure. It also stops the first notification from revealing the rest of a shared-vulnerability portfolio. If one flaw reaches many listed realms, the books close together and the operators are notified together. Some delay remains possible before the instruction reaches VMark. That delay exists in ordinary vulnerability research. The market is not a promise that every hole is disclosed at the moment it is found. It is a promise that, once the holder is ready to be paid, the books close and the operators are told together.

A proof that causes prohibited harm is invalid. A proof outside the listed boundary is invalid. A ransom note is not a proof. Exfiltrated customer data is not a proof. The market pays for the bounded marker and the capability-equivalent path, not for damage.

D. Verification

Verification has two layers. The marker establishes the outcome: the record was written, the file appeared, the canary was used, or the key signed. VMark then checks that the committed method caused that outcome and satisfied the contract. The first layer is often mechanical. The second is adjudication.

A lookup of _vmark does not establish that a DNS path was unauthorized. A callback does not establish remote execution. A valid signature does not establish that a key-use policy was bypassed. The disclosure, identity and access records, third-party telemetry, and reproduction or forensic evidence are what establish those facts.

VMark is the primary adjudicator. A listing may name an independent laboratory to assist. A financially interested operator cannot accept or reject a claim on its own. The listing fixes the verifier, the evidence required, the decision period, the conflicts policy, and the appeal authority before trading begins. Where remediation prevents literal reproduction, the proof class says what contemporaneous evidence is enough.

Participants have to trust VMark to keep disclosures, not trade on them, apply the published rules, and issue consistent decisions. VMark in turn has to be audited, to keep adjudication apart from trading, to protect the disclosure repository, and to maintain an appeal. The signed marker and the ruling make the observable parts independently checkable. They do not make private causation prove itself.

Collusion by an operator, a claimant, an employee, a verifier, or an affiliate is market manipulation. The marker format cannot eliminate it. Identity records, beneficial-owner records, position surveillance, preserved logs, conflicts rules, contractual penalties, and ordinary civil and criminal enforcement address it as they do other fabricated events.

E. Dynamics

The payout changes the probability of the settlement event. That is what reflexive means here. If the event were harm, the instrument would pay for causing harm. The event is instead a bounded demonstration and a private disclosure. The party best

able to predict settlement may be the party able to produce the proof. The cheapest qualifying way to force settlement should be the authorized marker, not the damaging action. If it is not, the proof class is wrong.

A lawful proof will not dominate every alternative at every pool size. It competes when the attainable net payout exceeds the value of continued secrecy, illicit sale, exploitation, and verification risk. VMark makes the lawful price visible and lets many sponsors put capital behind it. Exploitation remains possible, as it is today. The market does not have to abolish it in order to redirect a substantial share of activity. It has to pay.

Discovery time is private and cannot be verified. A holder may wait in the hope that L grows, or that more affected realms are listed. Waiting also risks another claimant closing the book, short dilution, remediation, or expiry. A holder who enters the short pool before submitting is still paid if someone else produces the proof, but the early position is a visible signal and may attract competing short capital. VMark does not infer when discovery occurred, and it does not run a coordinated research phase. It creates a race to establish and monetize an informed position against a finite pool. Whether this accelerates disclosure relative to the alternatives is a question for the settlement data.

After a tenor settles, the operator holds the proof, the shorts have been paid, and the operator remediates or does not. The next contract on the same realm prices the new state. If the condition is gone, parties who believe that go long. If the proof still reproduces, parties who can still produce it go short. The new $\hat{p}$ is an observable price after remediation. It is not a score of the patch. It also reflects funding, attention, liquidity, configuration changes, and other holes. No second contract on whether the operator will patch is required. An operator who wants a lower short ratio has to earn it on the next tenor.

F. Regulation

VMark is not a token and not a DAO. It is intended to operate through an institution that is permitted to list the contracts, hold customer funds, verify participants, surveil positions, adjudicate proofs, and enforce its rules. In the United States, event contracts may implicate the Commodity Exchange Act and CFTC rules, including review of contracts that involve or reference unlawful activity [6]. Recent CFTC guidance also emphasizes fraud and misuse of nonpublic information in event-contract markets [7].

A lawful proof holder must be able to take the proof side. That is not incidental. The trading rules have to distinguish lawfully acquired vulnerability knowledge from fabricated events, breaches of duty, operator self-dealing, verifier trading, and collusion. Identity verification, position reporting, restrictions on VMark and verifier personnel, special treatment of operator affiliates, and segregation of customer collateral are part of that form, not extras.

Listing is a legal instrument. The operator must have authority over the protected object and must specify the permitted conduct. The Department of Justice's CFAA charging policy recognizes the public value of good-faith security research designed to avoid harm, while making clear that the policy is not a private legal right and does not excuse bad-faith conduct [8]. VMark's contractual authorization, proof-class limits, identity records, and prompt private disclosure are intended to make that boundary clearer than it is in unsolicited research.

The first implementation should be narrow: volunteer realms, a small number of high-confidence proof classes, designated independent verifiers, bounded tenors, and capital supplied by operators or public sponsors. Market regulation is part of the construction. Its purpose is to make informed cyber trading legible and enforceable, not to eliminate informed participation.

References

[1] OpenAI, "A call for collective action on cyber defense," 27 August 2026. https://openai.com/collective-cyberdefense

[2] R. Böhme, "Vulnerability Markets: What is the Economic Value of a Zero-Day Exploit?" 22nd Chaos Communication Congress, 2005.

[3] O. Bohsali, "Reflexive Prediction Markets," July 2026. https://omarish.com/p/reflexive-prediction-markets

[4] J. Bell, "Assassination Politics," 1995.

[5] B. Franklin, *Poor Richard's Almanack*, 1735. "Three may keep a Secret, if two of them are dead."

[6] Commodity Futures Trading Commission, "Prediction Markets Advisory," CFTC Letter No. 26-08, 12 March 2026; 17 C.F.R. § 40.11. https://www.cftc.gov/csl/26-08/download

[7] Commodity Futures Trading Commission, "Advisory on Enforcement Authority over Event Contracts," 25 February 2026. https://www.cftc.gov/media/13351/Enf_AdvisoryKalshi022526/download

[8] U.S. Department of Justice, "9-48.000—Computer Fraud and Abuse Act," Justice Manual, updated 6 February 2025. https://www.justice.gov/jm/jm-9-48000-computer-fraud