Chronicle: Bringing Trust and Accountability to the Web

Luka Dover
https://chronicle-network.org

Abstract

The internet lacks a general public signal of trust. Existing signals such as likes, reviews, follows, and platform-local reputation are useful, but they are cheap to produce, easy to manipulate, fragmented across platforms, and often centrally controlled. We propose a protocol for costly public orientation on the web. Participants publish signed orientations toward or away from arbitrary subjects, expressed through publicly verifiable sacrifice. Nodes accept these events, relay them in real time, group them into batches, compute a Merkle root for each batch, and anchor that root to Bitcoin. After anchoring, nodes publish the full batch data off chain, allowing anyone to verify the anchor and reconstruct the public record. From this published data, downstream services can estimate trust profiles, derive subject scores, and compute actor reputation. Nodes do not need global coordination, and they do not need to provide full query interfaces. Their role is simply to accept valid events, anchor batch roots, and publish the corresponding batch data. We present the protocol structure and outline a baseline method for computing trust and reputation from the resulting public record.

1 Introduction

Interactions on the internet constantly depend on trust judgments. A user deciding whether to click a link, open an email, rely on a seller, follow an account, or use a service is always making some estimate of likely outcomes. Existing trust mechanisms such as reviews, ratings, follows, and platform-local reputation work well enough in many cases, but they remain cheap to produce, easy to manipulate, fragmented across platforms, and often centrally controlled. As a result, trust online remains local, weak, and easy to manipulate. A malicious website can appear legitimate, a spam sender can appear authentic, and a product can accumulate fake reviews that are difficult to distinguish from real ones. There is no widely adopted general-purpose internet protocol for signaling trust in a way that is portable, public, and accountable.

What is needed is a network for coordinating and aggregating distributed trust signals into a shared public map, allowing browsers, email clients, marketplaces, and social platforms to filter noise, surface substance, and respond more quickly to abuse. In this paper we propose a protocol that makes trust signaling costly, public, and portable. Instead of relying on cheap expression alone, it allows participants to publish signed orientation events toward or away from arbitrary subjects, backed by irreversible and provable sacrifice. Nodes accept these events, relay them, batch them, and anchor commitments to Bitcoin. From the resulting public data, downstream services can estimate trust profiles, derive subject scores, and compute actor reputation. The effect is to make manipulation more expensive, trust more portable, and collective response easier, so that public trust maps can emerge from the costly orientations of many participants.

2 Trust

Trust can be understood as a summary of anticipated outcomes. A subject is trusted when interaction with it is expected to lead mostly to positive outcomes, distrusted when negative outcomes predominate, and controversial when expectations are split.

Let V denote anticipated outcome value associated with a subject. We define trust as

$$\text{Trust} = \Pr(V > 0) - \Pr(V < 0) \in [-1, 1],$$

so that trust measures the balance of positive and negative anticipated outcomes.

A trust profile visualizes that distribution. In the simplified figure below, -1 and $+1$ denote negative and positive outcomes respectively.

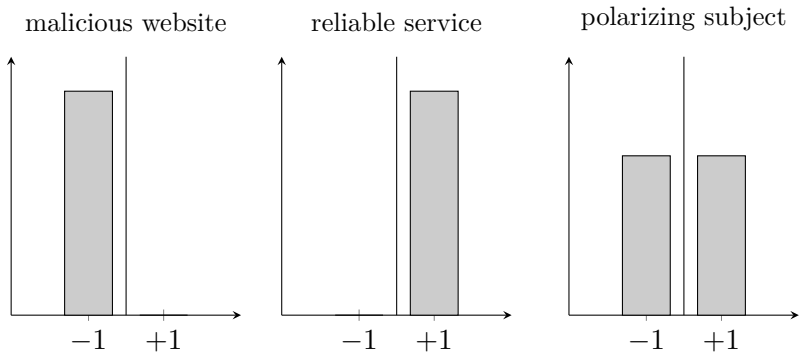

If the internet is to produce a public trust signal, it must produce credible evidence about that distribution. Existing systems do provide reviews, ratings, and warnings, but because those signals are cheap to produce and easy to manipulate, they cannot by themselves generate a reliable public map.

3 Costly Orientation Events

The solution starts with the concept of a costly orientation event: a signed signal toward or away from a subject that is expressed through irreversible sacrifice. The sign expresses direction and the amount expresses magnitude. The sacrifice is not a payment to another party, but a provable destruction of value. That distinction matters: payments can be routed, refunded, or cycled among cooperating actors, whereas destruction makes commitment publicly legible and difficult to game.

$$\texttt{event} = [\texttt{subject}, \texttt{signed\_amount}, \texttt{pubkey}, \texttt{timestamp}, \texttt{sig}]$$

A positive amount orients toward the subject. A negative amount orients away from it. The signature proves that the event was created by the holder of the corresponding public key.

For example:

```
["scam.shop", -100, user_pubkey, timestamp, sig]
["example.com", +50, user_pubkey, timestamp, sig]
```

A subject may be any protocol-defined identifier or entity. Subjects can include domains, URLs, pubkeys, nodes, user accounts, or prior orientation events. This allows the same primitive to represent first-order trust judgments, direct judgments about actors or infrastructure, and later acknowledgment of whether a prior orientation proved useful.

4 Nodes

Costly orientation events need an execution layer, which is provided by public nodes. A node is a paid public service that accepts valid signed events from clients, charges a fee for processing them, groups accepted events into a batch, computes a cryptographic root for that batch, and anchors that root to Bitcoin [1]. The anchor transaction must destroy at least as much bitcoin as the total sacrifice represented by the events in the batch. This allows many orientation events to share a single public anchor while preserving independent verifiability.

For each batch, the node orders accepted events deterministically, hashes them into a Merkle tree [3], binds the resulting root to its own public key, and places that commitment in an `OP_RETURN` output of a Bitcoin transaction.

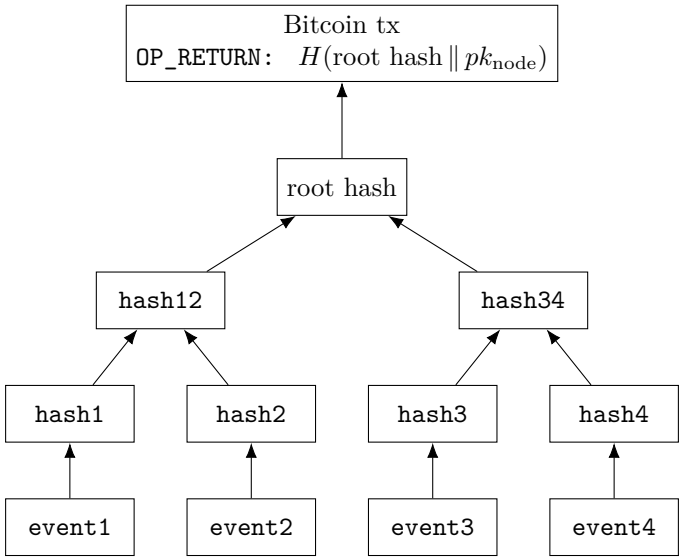

Events are grouped into a batch, hashed into a Merkle root, and anchored to Bitcoin.

After anchoring, the node publishes the full batch data off chain. Anyone can then recompute the batch root from the published events and verify that the corresponding Bitcoin transaction contains the correct commitment and sufficient sacrifice. Published batches therefore form a node's public operational history, which anyone can inspect to evaluate performance and reliability.

5 Real-Time Coordination

Bitcoin anchoring provides public finality, but many coordination problems require a faster response than block times allow. A scam link, malicious sender, or compromised account may need to be signaled immediately, not only after the next anchor is published. To support real-time coordination, nodes relay accepted events as soon as they are received.

When a node accepts an orientation event, it immediately broadcasts the event together with a receipt. The event shows that the actor signed the orientation. The receipt shows that the node accepted responsibility for including the event in a later anchored batch. This allows clients and downstream services to respond before final anchoring on Bitcoin.

Receipts make node responsibility public. If a node repeatedly emits receipts that are not later honored in published batches, that failure becomes visible in the public record. The node can then itself become the subject of negative orientation within the same protocol, damaging its standing

and reducing future trust in its service.

6 Incentives

The protocol relies on incentives at three levels: the incentives of nodes that execute and publish events, the incentives of participants whose orientations shape the trust map, and the incentives of downstream service providers that store, index, audit, and serve useful views of the public data.

Nodes are paid public services. They charge fees for accepting events, relaying them, anchoring them, and publishing verifiable batch data. This gives nodes a direct incentive to remain available, honor receipts, and maintain a public record of reliable work. A node that censors events, fails to publish its batches, or repeatedly emits receipts that are not later honored, damages its own standing and risks losing future traffic to competing nodes.

Participants are incentivized not only by direct influence over trust profiles, but also by reputation. An orientation that later proves useful for others can itself become the subject of later positive orientation. In this way, actors who repeatedly help others avoid harm or find value can accumulate public reputation and with it greater influence over future trust-map computation.

Nodes are execution and publication services, not query engines. Downstream service providers can build on top of the published data to provide storage, indexing, trust-map APIs, reputation computation, and node auditing. These services can also be monetized. Providers that make the public data easier to access, interpret, and use have a direct economic incentive to improve the quality of the ecosystem while helping keep nodes accountable. As more services mirror and process published batches, storage becomes more redundant and the public record more resilient.

The incentives of the system are therefore recursive. Nodes are rewarded for operational honesty and reliability. Participants are rewarded for producing orientations that others later find useful. Downstream service providers are rewarded for making the public trust record usable. Manipulation is not eliminated, but it becomes more expensive, more public, and more accountable.

7 Computing Trust

Services that compute trust lie outside the strict protocol boundary. The protocol itself publishes a public graph of costly orientations; downstream services interpret that graph and derive useful signals from it. Different services may use different algorithms, and those algorithms will evolve over time. Their design nonetheless matters, because it shapes how the protocol is used, how influence accumulates, and how the public record is interpreted. For example, an algorithm that rewards orientations toward prior events creates an incentive to publish explicit usefulness judgments about earlier signals. For that reason we outline a baseline trust-computation algorithm.

Published orientation data is useful even in the simplest case. A first estimate of trust can be obtained by aggregating net signed sacrifice toward a subject. If many participants independently orient away from a domain, the resulting trust profile shifts negative. If many orient toward it, the profile shifts positive. Even this simple approach already produces a stronger public signal than systems based on costless ratings or reviews.

A better algorithm, however, should not treat every source equally. Some actors prove more useful than others. An orientation made by a participant whose prior signals repeatedly helped others avoid harm or find value should count more than an orientation made by an unknown or consistently misleading actor. Raw trust profiles should therefore be weighted by reputation. Wealth can buy signal magnitude, but it should not automatically buy influence over the map, because influence should depend on an actor's public record of producing useful orientations.

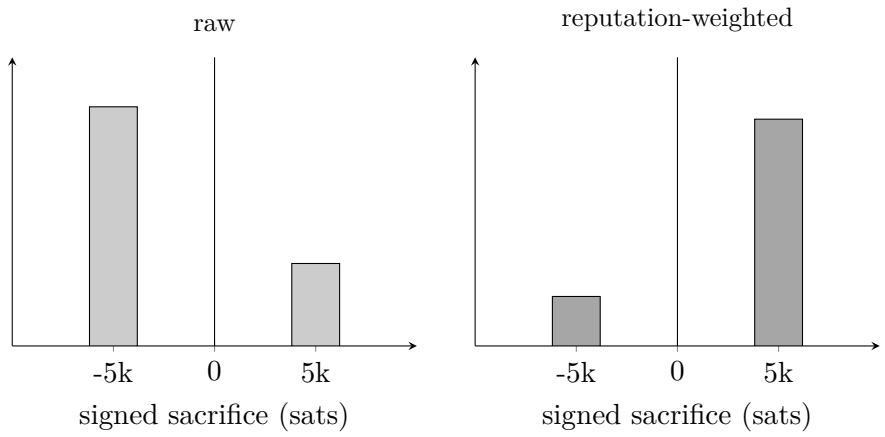

The same raw sacrifice can imply a different trust profile once source reputation is taken into account.

This makes trust computation recursive. A subject's score depends on the costly orientations pointed toward it, weighted by the reputations of the actors who made them. An actor's reputation depends in part on whether that actor's prior orientations later proved useful. That usefulness becomes visible because orientation events can themselves become subjects of later orientation. A participant who benefits from a prior signal can orient toward that specific event, making its usefulness explicit rather than leaving it to be inferred indirectly.

Our baseline proposal computes three linked quantities: a subject score Q, an actor reputation R, and an event usefulness score U. The full update equations are given in the appendix.

The intuition is similar to PageRank [2]. PageRank showed that importance can be extracted recursively from a web of links. Here the same general idea is applied to a different object: a public graph of signed costly orientations. Orientation events are explicit, directional, and backed by sacrifice. They can point toward or away from arbitrary subjects such as domains, accounts, public keys, email addresses, phone numbers, or prior orientation events.

That contrast matters. Early Google succeeded because the web's link structure still carried traces of real cost. Creating and curating links required time, judgment, and attention, so links functioned as a rough proxy for orientation. But because that signal was only inferred from a proxy, it could eventually be gamed. Once the metric became legible and profitable to manipulate, the proxy detached from the underlying cost and the ranking signal degraded, a familiar instance of Goodhart's Law [4].

The aim here is similar recursive extraction, but from a graph whose magnitudes are explicit, signed, and backed by real cost.

8 Resistance to Manipulation

A public trust map becomes harder to bend as costly history accumulates. An actor can always spend money and publish new orientations, but influence should not come cheaply. Reputation must be earned through a visible record of orientations that others later find useful. That takes time, repeated sacrifice, and later acknowledgment from participants who benefited from those signals. Once earned, reputation gives an actor greater weight in future trust computation. But using that influence to push harmful or misleading orientations should place the same reputation at risk.

The protocol provides the public graph of orientations, but trust-computation algorithms determine how that graph is interpreted. Those algorithms therefore shape behavior: they influence

which actions build credibility, how quickly new sacrifice can move the map, and how easily accumulated influence can be abused. To resist manipulation, robust algorithms should account for three structural properties.

Capital inertia. A graph with a large body of anchored sacrifice should be hard to move. In a mature graph with millions of participants and a large accumulated history of sacrifice, a new manipulator should need substantial sacrifice before the map shifts meaningfully.

Historical advantage. A long public history of sacrifice should count more than an equal amount that appears suddenly. Time matters because a long public record has had more opportunity to be exposed to correction, contradiction, and acknowledgment.

Reputational fragility. Reputation should be slow to earn and easier to lose. An actor may gain influence through a sustained public record of useful orientations, yet lose that influence quickly once later evidence reveals a pattern of harmful or misleading behavior.

These are not properties of the protocol, but of the algorithms built on top of it. The protocol makes these properties computationally available by publishing a public graph of signed, costly, time-ordered orientations that can target subjects, actors, and prior events.

9 Conclusion

We have proposed a base-layer protocol for public trust signaling on the internet. Instead of relying on cheap, fragmented, and platform-local signals, the protocol makes orientation costly and publicly verifiable. Nodes operate without global coordination. They accept fee-paying events, relay them, group them into batches, anchor batch roots to Bitcoin, and publish the corresponding batch data off chain. Downstream services can then use that public record to estimate trust profiles, compute reputation, and build useful interfaces on top.

We have also outlined a baseline approach to trust computation in which subject scores are shaped not only by sacrifice, but by the public usefulness of their sources. In this way, the same orientation graph that records costly action can also support recursive reputation and increasing resistance to manipulation over time.

The protocol is not a replacement for reviews, ratings, or existing reputation systems. Reviews provide context and explanation; Chronicle adds magnitude and accountability through publicly verifiable sacrifice.

The protocol changes the economics of influence by making trust signaling costly and accountable. It offers a shared coordination layer through which trust-relevant judgments can become more public, more reusable across contexts, and more resistant to manipulation across the web.

A Baseline Q/R/U Update Equations

The baseline proposal computes three linked quantities iteratively: subject score Q, actor reputation R, and event usefulness U.

Let a_e denote the signed sacrifice attached to event e. Let actor(e) denote the public key that authored event e. Let $e \to s$ mean that event e points to subject s, $e' \to e$ that event e' points to prior event e, and $e \to p$ that event e points directly to actor p. Let E_p be the set of events authored by actor p.

$$Q_s^{(t+1)} = \sum_{e \to s} a_e \, R^{(t)}_{\mathrm{actor}(e)}$$

$$U_e^{(t+1)} = \sum_{e' \to e} a_{e'} \, R^{(t)}_{\mathrm{actor}(e')}$$

$$R_p^{(t+1)} = \alpha \sum_{e \in E_p} U_e^{(t)} + \beta \sum_{e \to p} a_e \, R^{(t)}_{\mathrm{actor}(e)}$$

In words: a subject's score is the reputation-weighted sum of orientations toward it; an event's usefulness is the reputation-weighted sum of later orientations toward that event; and an actor's reputation is updated from both the usefulness of that actor's prior events and direct orientations toward the actor.

This is a minimal baseline rather than a finished standard. Different downstream services may choose different weighting, normalization, damping, or time-decay schemes.

References

[1] Satoshi Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. 2008. https://bitcoin.org/bitcoin.pdf

[2] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. *The PageRank Citation Ranking: Bringing Order to the Web*. Stanford InfoLab, Technical Report 1999-66, 1999.

[3] Ralph C. Merkle. A Digital Signature Based on a Conventional Encryption Function. In *Advances in Cryptology — CRYPTO '87*, Lecture Notes in Computer Science, Vol. 293. Springer, 1988.

[4] Charles A. E. Goodhart. Problems of Monetary Management: The U.K. Experience. In *Papers in Monetary Economics*, Vol. 1. Reserve Bank of Australia, 1975.