

Sentinel Signup

Registration screening for approval queues

Sentinel Signup

Registration screening that scores behavior, never newness

Open Feed Network, Inc. · candortrustandsafety.com *Every claim below maps to a dated behavioral proof in the public evidence record.*

The problem

Open registrations invite bot farms; closed registrations kill communities. Most instances resolve this by making moderators read every application — which works until the day a farm files fifty.

What Sentinel Signup does

It screens the approval queue automatically so moderators only see the cases that genuinely need human judgment. Bots are rejected at the edge, humans are approved in seconds, and ambiguous cases are left where they belong — with your moderators.

How it works, in principle

Applications are scored on behavior, never on newness. Being a new account is not a signal; acting like a coordinated batch is — arrivals clustered in time, “different” applicants sharing network space, and identical or generated application language across accounts. Verdicts fall into three bands: pass (reads human, approved without a moderator touching it), flag (ambiguous, held for human review — the moderator decides, never the machine), and block (a coordinated pattern established, rejected at the edge). *The specific signals and scoring are not published.*

Verified behaviors

- **Coordinated subnet signups are flagged, then blocked as the pattern accumulates.** *Verified July 27, 2026.*
- **Legitimate human applications pass in seconds, before and after an attack.** *Verified July 27, 2026.*

- **Verdicts survive service restarts.** *Verified July 27, 2026.*

Honest limitations

Language or locale never auto-rejects an application — only behavior does. Blocks expire by design; the system biases toward letting humans back in. This screens the front door only — it is registration screening, not content moderation after admission.