

A PUBLIC-INTEREST ACCOUNTABILITY NETWORK

# Integrity *Alliance*

*Full product briefs — investigation, verification, protection, and vetting, built as one connected infrastructure for institutions that interact with children, survivors, and people in crisis.*

---

<b>Founded</b>	2026
<b>Form</b>	501(c)(3) Public Charity in formation
<b>Editorial</b>	Independent
<b>Method</b>	Primary-source

---

## I N T H I S D O C U M E N T

- I. One-Pager — Founding Partner Introduction
- II. Discern — Product Brief
- III. Ambience Integrity — Product Brief
- IV. Founding Pitch — February 2026

# I.

I N T E G R I T Y   A L L I A N C E

## One-Pager

*We hold the institutions closest to the vulnerable.*

*A 501(c)(3) accountability network.*

February 2026 · A founding partner introduction  
founding@integrityalliance.org · Watchdog Report Vol 1, Issue 1 available on request

---

### Who We Are

Integrity Alliance is an independent accountability organization investigating the corporate, financial, and institutional structures that hold power over vulnerable populations — children, agricultural and residential-care workers, crisis-line callers, victims of violence — wherever private power has unaccountable access. We are not a think tank. We publish.

### What We Do

The Watchdog Report. Quarterly investigative journalism into ownership chains, regulatory records, and corporate custody of vulnerable people’s data and bodies.

Discern. Verification of synthetic media, AI personas, and impersonation in the age of generative AI. Engine: TruthMark Verify.

Sentry Standard. Independent vetting of adults, facilities, and AI personas with access to children — applied identically across public, private, and government entities.

### Proof of Concept

Volume 1, Issue 1 of the Watchdog Report shipped in February 2026. Two investigations, twenty-plus primary sources (DOJ, EPA, SEC, Senate Finance Committee, IARC, federal court filings), a four-persona editorial audit completed before release, editorial disclosures at the top of each story, and a published revision log. The institution we are proposing to scale already produces the work it exists to scale.

### What We’re Asking

Founding Circle. Three to five founding gifts of \$100K–\$250K to seat Year 1 operations.

Advisory Council. Members across child protection, investigative journalism, First Amendment law, OSINT, AI integrity, clinical practice, and faith-community leadership.

Editorial Board. Five to seven independent governors with staggered three-year terms; faith-community and secular representation by design.

Operating budget Year 1. \$500K–\$750K.

Protect the vulnerable. Document the truth. Publish the corrections.

# II.

P R O D U C T   B R I E F   ·   D I S C E R N

## Discern

*In the Age of AI and Phantom Realities.*

v1.0 | February 2026 | A product of Integrity Alliance | Engine: TruthMark Verify

Discern — Product Brief, Version 1.0 | A Integrity Alliance Product | Engine: TruthMark Verify

---

### What Discern Is

Discern is the verification layer for an internet that no longer reliably tells the truth about itself. Built on TruthMark Verify’s existing engine and repositioned for child-safety, civic-integrity, and parental audiences, Discern is the public-facing surface where a person can paste a username, an image, a video, or a recording and receive a structured account of what the content has the technical signals of, what it lacks, and where verification confidence breaks down.

Discern is not a fact-checker. It does not adjudicate truth. It surfaces evidence and names confidence. The decision is the user’s.

### The Problem

A child befriends an account online. A parent sees a viral clip of a public figure saying something extraordinary. A journalist receives a leaked recording. None of them can independently verify whether they are looking at a person, a clone, a fully synthesized persona, or one cloned voice reshared across six channels. The default trust posture of the platforms is to assume authenticity. Discern reverses that default.

### Core Capabilities

Discern delivers verification across four signal areas:

**Identity provenance** — Username and persona search across major platforms; cross-platform identity correlation; account-age and behavioral-pattern signals.

**Visual integrity** — Reverse image search; deepfake artifact scoring; C2PA / content-credentials parsing where present; metadata and EXIF inspection.

**Audio integrity** — Voice-clone detection; cross-recording pattern matching (is this the same cloned voice across multiple channels?); synthetic-speech artifact scoring.

**Source provenance** — Original-content tracing; cross-platform syndication mapping; the chain from first appearance to viral moment.

## Methodology

Discern inherits its editorial discipline from the Watchdog Report standard:

Multi-tool corroboration: no single tool's signal is dispositive.

Confidence-level disclosure: every Discern report names its certainty band — high-confidence synthetic, likely synthetic, inconclusive, likely authentic, high-confidence authentic.

Tools are advisory, not adjudicative. Discern surfaces evidence; humans decide.

All published Discern Reports meet the same primary-source standard as the Watchdog Report and pass the same four-persona editorial audit before release.

## Product Surfaces

**Verify Tools** — Free public access. The user submits a username, image, video, or audio file and receives a structured verification report. Hosted on the TruthMark Verify infrastructure; surfaced under the Discern brand for child-safety and civic audiences.

**Discern Reports** — Editorial investigations into specific synthetic-media incidents, published as features in the Watchdog Report. Founding investigation: the cloned-voice protest content distributed across six YouTube channels, TikTok, and broadcast television, cross-corroborated with parent reports of school-day protest recruitment of minors.

**Discern Alerts** — Subscription notifications when emerging synthetic-media patterns reach children's spaces — impersonator accounts on gaming platforms, AI-generated influencer personas targeting minors, voice-clone trends in family-facing media.

## Privacy and Data Posture

Discern processes content provided by the user. No identifying information about the subjects of verification is retained beyond the active session unless the user explicitly archives a report. Where children's data appears in Discern's published investigations, handling is governed by the forthcoming children's-data tokenization standard (Sentinel).

No content uploaded to Discern is sold, shared with third parties, or used to train commercial AI models. The verification engine runs against signals — not the original media — after analysis.

## Limits — Honestly Named

Discern cannot prove a video is real. It can score the likelihood of synthetic manipulation.

Discern cannot identify the creator of a deepfake. It can demonstrate that one was likely used.

Detection capability degrades against state-of-the-art generation models. We update continuously and publicly disclose model freshness.

Discern does not replace human judgment. It augments it.

## **Relationship to the Network**

Discern is one of three operational arms inside Integrity Alliance — alongside the Watchdog Report (flagship investigative publication) and the proposed Sentry Standard (vetting layer for entities and personas seeking access to children).

Discern provides the verification engine. The Watchdog Report provides the editorial discipline and platform of record. Sentry Standard applies vetting where children are involved. Discern is the bridge between technical verification and journalistic publication: every Discern Report ultimately ships under the Watchdog Report standard.

# III.

PRODUCT BRIEF · AMBIENCE INTEGRITY

## Ambience Integrity

*Forensic environmental psychology — waterfall wellbeing through every business atmosphere.*

v1.0 | February 2026 | Sister product of Integrity Alliance | Live at [ambience-integrity.vercel.app](https://ambience-integrity.vercel.app)

Ambience Integrity — Product Brief v1.0 | Sister Product of Integrity Alliance | Live at [ambience-integrity.vercel.app](https://ambience-integrity.vercel.app)

---

### What Ambience Integrity Is

Ambience Integrity is a forensic environmental psychology platform that measures how color, light, sound, and layout flow through any physical business space. The framework applies universally — schools and playgrounds, daycares and clinics, factories and warehouses, restaurants and retail, offices, event venues, residential care facilities, places of worship, transit hubs, and every other built environment where humans gather. Operators use it to score and remediate the wellbeing of their environments across four pillars: safety, health, welcoming atmosphere, and revenue.

The same methodology serves a second purpose. Forensic environmental psychology applied with discipline reads what a corporate facade conceals. A facility that violates the framework on every measurable dimension cannot hide that record from a trained reviewer. Ambience Integrity is therefore both an operator self-improvement tool and an investigative methodology — the field framework behind any Integrity Alliance investigation that requires structured assessment of a physical space.

### The Four Pillars

Every Ambience Integrity audit evaluates a space across four documented dimensions:

**Environmental safety.** Lighting, contrast, noise, egress, materials. The variables that predict slips, collisions, hearing loss, supervision failures, and physical risk to occupants. Regulators and insurers score what they can document; unseen hazards become citations, claims, and shutdown risk.

**Worker and visitor health.** Glare, VOC-heavy materials, thermal stress, acoustic exhaustion. The chronic variables that drive absenteeism, errors, and stress — and the environmental sources of OSHA-recordable patterns that get falsely attributed to behavior.

**Consumer and client welcoming.** Wait perception, color arousal, wayfinding, scent policy. The variables that shape trust before a person speaks to staff. Customers and patients who feel safe and calm stay, return, and review well.

Revenue maximization. Dwell time, conversion paths, consultation closure, table turn. The variables that owners optimizing only spreadsheets overlook — the highest-leverage operational lever most operators do not control deliberately.

## Industry Verticals — Shipping and In Development

The framework applies to every industry. The platform currently ships with documented reference atmospheres, approved color palettes, and verticalized playbooks for nine industries, with additional verticals in continuous development. The shipping nine:

**Event planner office** — serene luxury, sage consultation suites, portfolio sightlines.

**School playground** — grounded play, fall-zone depth, shade on metal slides, supervisor sightlines.

**Daycare** — nurturing energy, child-scale furniture, saturated-accent limits to prevent overstimulation.

**School** — clear learning, 300–500 lux on work surfaces, glare-free whiteboards, illuminated egress.

**Restaurant** — hospitality glow, warm dim evening light, host-stand de-escalation tones.

**Medical and dental** — clinical calm, blue-green wait areas, visible infection-control cues.

**Retail** — inviting flow, conversion-path lighting, fitting-room comfort, returns-desk de-escalation.

**Factory and industrial** — hazard-coded color, hearing-conservation zones, recovery rooms under 55 dB.

**General office** — professional calm, glare-free screens, acoustic pods, biophilic break zones.

New verticals are added as operators in those industries request playbooks and as Integrity Alliance investigations require structured assessment frameworks for additional space types. The four pillars and the compliance ladder remain constant; the verticalized playbooks adapt the framework to the specific safety, health, welcoming, and revenue variables of each industry.

## The Compliance Ladder

Each audit places a space on a four-tier compliance journey, with named remediation actions at every tier:

**Non-compliant** — Critical playbook failures, open citations, or repeated complaints. Stop-the-bleed fixes within 48 hours; insurer and legal notification where regulatory notice has been received.

**Working toward compliant** — Remediation in progress with a documented 90-day roadmap, weekly photo evidence on critical rules, scheduled enforcement check-ins.

**Compliant** — No critical open findings; warnings have scheduled fixes; meets ADA/OSHA-aligned playbook. Re-audit at 6–12 months.

**Fully compliant** — Verified remediation, stable score, optional carrier-discount tier, BBB- and government-response-ready evidence pack.

## Operator Tools

**Learn** — Psychology, compliance, and hazard guides organized by vertical. The entry point for first-time operators.

**New Audit** — Score a real location against the playbook for its vertical.

**GFly Assistant** — On-site copilot. The AF GFly local research assistant integrated as field-mode companion during audits.

**Whiteboard Academy** — Operator training through The Whiteboard adaptive tutor system. Builds environmental literacy across staff, not just owners.

## **Dual Use — Operator and Investigator**

The same forensic environmental framework that helps a daycare operator improve their classroom can document a Bonnie Plants facility from the outside in. The discipline of the four pillars — safety, health, welcoming, revenue — produces structured evidence whether the goal is remediation or investigation. Whistleblower observations gain rigor when categorized inside the framework. Physical-space anomalies that would otherwise be dismissed as narrative become evidence categories that survive scrutiny.

This dual use is intentional. Most operators are not bad actors; most need the framework to improve. The minority of operators who are bad actors cannot disguise sustained framework violations across all four pillars. Ambience Integrity serves both, transparently, with the same methodology.

## **Relationship to the Alliance**

Ambience Integrity is a sister product of Integrity Alliance rather than an internal arm. Its commercial use — operators paying for self-audits and remediation guidance — funds its continued development independently of Integrity Alliance grant revenue. Its methodology is shared open-source so Watchdog Report investigations can use it without licensing dependency.

Where Ambience Integrity audits surface findings that meet Integrity Alliance investigative thresholds, those findings can be routed through Watchdog Report editorial channels with operator and subject consent. Where Sentinel applies, any individual data in Ambience Integrity audits affecting protected populations is tokenized per the Sentinel standard.

# *IV.*

FOUNDING PITCH · FEBRUARY 2026

## **The Founding Pitch Deck**

*The original ten-slide founding pitch follows on the next pages, reproduced slide-for-slide from the source deck — layout, palette, and copy intact.*

---

FOUNDING PITCH · FEBRUARY 2026

# Integrity Alliance

---

Independent Investigation. Verified Reporting. Protected Children.

Confidential — Founding Document Series

WHY WE EXIST

# We didn't start with a pitch deck. We started with an investigation.

The Watchdog Report — Volume 1, Issue 1 —  
published in February 2026. Two investigations.  
Primary sources. Editorial disclosures at the top of  
each story. A published revision log. A  
four-persona editorial audit applied before release.

The institution we are proposing already produces  
the work it exists to scale.

WATCHDOG REPORT · VOL 1, ISSUE 1

**2**  
investigations published

**20+**  
primary sources cited

**4**  
editorial-audit personas

## THE REAL PROBLEM

**Two converging crises. No independent third party watching.**

**01**

**Private equity now owns the companies holding our children's most personal data.**

Lifetouch → Shutterfly (\$825M, 2018) → Apollo Global Management (\$2.7B, 2019) → Leon Black, who paid Jeffrey Epstein up to \$170M between 2012–2017. No parent ever received a letter.

**02**

**AI-generated voices and synthetic personas are infiltrating spaces meant for kids.**

A single cloned voice surfacing across six YouTube channels, TikTok, and broadcast television — mobilizing school-day protest recruitment of minors. No regulator owns this beat.

In both cases, the institutions closest to our children are governed by entities our children's parents have never met.

WHAT WE DO

# Investigate. Verify. Report.



## Investigate

OSINT collection on ownership chains, regulatory records, financial flows, and data practices of entities that interact with children, vulnerable populations, and people in crisis.



## Verify

Two independent confirmations per factual claim. Four-persona editorial audit — forensic investigator, senior editor, child protection expert, military specialist — before any story ships.



## Report

Chicago-style endnotes. Editorial disclosures front-of-piece. Published revision logs. Internal taxonomy of Verified / Caution / Avoid — with no letter grades on crisis services.

**OUR PLEDGE: We do not publish letter grades on crisis services. A poorly framed grade on a hotline can deter the very people who need to call.**

## THE NETWORK

# One parent organization. Three focused arms.

## Integrity Alliance

501(c)(3) parent · governance,  
methodology, legal

### The Watchdog Report

#### FLAGSHIP INVESTIGATIVE JOURNALISM

Quarterly investigations into ownership chains, regulatory records, and corporate custody of children's data. Already publishing.

### Discern

#### AI · SYNTHETIC MEDIA INTEGRITY

“Discern in the Age of AI and Phantom Realities.” Voice cloning, AI personas, and synthetic content that targets or impersonates real people.

### Sentry Standard

#### PROPOSED · CHILD-FACING VETTING LAYER

Independent vetting of facilities, adults, and AI personas with access to children. Working name; final naming under review.

Faith-based financial transparency, if pursued, will be incubated as a separate independent organization — not under the Integrity Alliance.

PROOF OF CONCEPT

# What we already do.

THE WATCHDOG REPORT

## Volume 1, Issue 1

Published February 13, 2026

**STORY ONE** · Bonnie Plants, Scotts Miracle-Gro, and the corporate ecosystem behind the 3rd Grade Cabbage Program.

**STORY TWO** · Lifetouch → Shutterfly → Apollo → Leon Black → the Epstein files.  
What the structural chain actually shows.

### METHODOLOGY APPLIED

Primary sources only

DOJ, EPA, SEC, Senate Finance, IARC, Bayer disclosures, Beyond Pesticides filings, corporate press releases

Editorial disclosures front-of-piece

Both stories open with explicit statements of what we are not alleging

Four-persona editorial audit

Forensic investigator, senior editor, child-protection expert, military specialist

Full public revision log

Every correction from v1 to revised edition documented and published

## 18-MONTH ROADMAP

# Sequenced for credibility before scale.

### 1 MONTHS 1–3 Publish

Methodology page live. Newsletter relaunch with subscriber list. Watchdog Issue 2: the synthetic-voice protest investigation

### 2 MONTHS 4–6 Staff

Hire Executive Director and lead investigative researcher. Secure \$500K–\$750K founding funding. Editorial Board seats first three members.

### 3 MONTHS 7–12 Scale

Issues 3 and 4 ship. Methodology submitted for peer review by journalism-school faculty. Subscriber growth to 5,000–10,000 engaged readers

### 4 MONTHS 13–18 Launch

Discern and Sentry Standard incubated. Secure \$1.5M–\$2M operating funding. Begin Congressional and state-AG engagement

**EVERY MILESTONE IS A PUBLICATION. We do not raise on a roadmap. We raise on what readers have already received.**

## REVENUE MODEL

# Diversified. Aligned. Independent.

**YEAR 1**

**\$500K–\$750K**

**YEAR 2**

**\$1.2M–\$1.5M**

---

<b>Founding Circle gifts</b>	Three to five \$100K–\$250K founding donors	<b>35%</b>
<b>Newsletter subscribers</b>	Individual \$50–\$150/yr; institutional \$500/yr	<b>20%</b>
<b>Aligned foundations</b>	Bradley, Searle, Marble Freedom Trust, journalism funders	<b>25%</b>
<b>Individual major donors</b>	Child-safety, civil-liberties, and accountability donors	<b>15%</b>
<b>Earned and in-kind</b>	Pro bono First Amendment counsel; methodology licensing	<b>5%</b>

## RISK & MITIGATION

**Naming what could go wrong — and how we've already designed for it.**

<b>Defamation</b>	<b>HIGH</b>	Primary-source-only standard. Two-source verification. \$2M+ media liability insurance. Pro bono First Amendment counsel pre-publication.
<b>Political backlash</b>	<b>HIGH</b>	Same standards applied evenly to faith-based and secular subjects. Public methodology. Diverse Editorial Board. Firewall from any advocacy work.
<b>Funder pressure</b>	<b>MED</b>	Editorial Independence Charter (legally binding). No single source >25%. Quarterly funder disclosure.
<b>Reporting errors</b>	<b>MED</b>	Independent fact-check pass. Four-persona editorial audit. Published revision log on every issue.
<b>Data breach</b>	<b>LOW</b>	Tokenization layer for any data on children. SOC 2 controls. Annual penetration testing. Minimal PII retention.

## THE ASK

# Join us as a founding partner.

---

### FOUNDING CIRCLE

Three to five founding gifts at \$100K–\$250K to seat Year 1 operations.

### ADVISORY COUNCIL

Twelve to fifteen members across child protection, investigative journalism, First Amendment law, OSINT, AI integrity, clinical practice, and faith-community leadership.

### EDITORIAL BOARD

Five to seven independent governors with staggered three-year terms; faith-community and secular representation by design.

### OPERATING BUDGET

\$500K–\$750K Year 1 to support Executive Director, lead researcher, legal counsel, and quarterly publication.

Protect the vulnerable. Document the truth. Publish the corrections.