

ce

AdCP Governance Areas

Launched 11/7/25

—

GROUP MISSION STATEMENT

Governance in and of itself is not an Agent, the work contributes to all Agents that are building AdCP. Consider which areas are: Industry thought leadership pieces vs built into the protocol

Define and operate transparent, privacy-safe, brand-suitable campaign governance that encodes human heuristics in open protocol objects, making campaigns easy to launch, safe to scale, and measurable against verifiable outcomes for buyers, sellers, and auditors.

1. Brand Safety & Suitability* bring in DV, IAS etc. (pause until further word from Brian per Lauren)

Group Working Doc -  **Data Privacy Laws and Regulations in Digital Advertising**

 **AdCP Brand Safety and Suitability Framework**

Definition:

Establish standardized, transparent safety parameters to ensure ads run in safe, appropriate, and compliant contexts.

Key Themes:

- Tiered safety rules: non-negotiable vs. flexible
- Embed safety logic into agentic workflows
- Balance automation with accountability and explainability

Open Questions:

- Who certifies AI-generated content for brand safety?
 - The entity **controlling** the ad creative is the primary party responsible
 - The entity **serving** the ad creative is a separate party acting as a processor under data protection concepts
 - The entity **displaying** the ad creative on their property is the second party responsible
 - The entity **classifying** the ad creative is a separate party acting as a processor under data protection concepts
- Should enforcement live in buy-side, sell-side, or a neutral governance agent?

A big governance danger: vendors trying to shove their flavor of brand safety into the object. Should the protocol enforce taxonomy-agnostic structure? Should vendors submit scores via namespaced fields? Should there be a “baseline” classification system + optional extensions?

Why it Matters:

It's the trust foundation of the system. If AI-driven buying can't ensure safe placement, adoption stalls. Governance must codify what “safe” means and how that translates into agent behavior.

Overlap:

Intersects heavily with **Creative Standards** (what's inside the ad) and **Measurement & Verification** (proving safety compliance). Relies on **Process & Oversight** to interpret gray areas that automation can't.

Governance Architecture requires a unified, taxonomy-agnostic safety object that every agent can reference, enforce, and log consistently across platforms.

Sub-Working Group members:

Lauren Carbone - Lead

Leo Boukou-Poba - Scope3

Mathieu Desvé

Tina Iannacchino - Seedtag

Jason White - Mula

Charley Saltzgaber - Triton Digital (c.f. [Sunder.ai](#))

2. Brand Manifesto & Standards Library* consider independent piece vs built in

Group Working Doc - [AdCP Brand Manifesto & Standards Library Framework](#)

Definition:

Create a machine-readable, privacy-safe library of brand values, exclusions, and approved vendors.

Key Themes:

- Unify scattered policies into a structured schema

- Allow brand-specific data-sharing permissions
- Tokenize sensitive elements for security

Open Questions:

- How do we balance interoperability with data confidentiality?
 - **Interoperability** should be considered sending / receiving unaggregated information in real-time across systems
 - **Confidentiality** should be associated with the contents of the information communicated AND use rights (which probably needs a separate noun like **Permissions/Restrictions**)
 - The concern of keeping business trading partners “confidential” conflicts with open systems design
 - The concern of keeping WHAT is being sold “confidential” can be achieved via look-up tables or encryption
- Should the schema live inside AdCP or as a linked sub-protocol?

Why it Matters:

This is the “memory layer.” It stores every brand’s ethics, exclusions, and tone of voice so agents know what’s acceptable. Without it, automation repeats mistakes at scale.

Overlap:p

Connects directly with **Brand Safety** (execution of rules), **Compliance** (legal alignment), and **Governance Architecture** (how the library is accessed and authenticated).

Governance Architecture needs a secure, versioned manifest schema that agents can lookup, authenticate against, and apply uniformly in all workflows.

Sub-Working Group members:

Amy Good - Strategy, Lead

Lauren Carbone

Travis Lusk

Leo Boukou-Poba

Pia Malovrh - Celtra

Stu Richards

Simon Safhalter - Celtra

3. Creative Standards & Validation* ask companies do they want a standard (Boltive)

Group Working Doc - [AdCP Creative Standards & Validation Framework](#)

Definition:

Ensure creative assets meet technical and ethical criteria before activation.

Key Themes:

- Integrate creative validation pre- and in-flight
- Align with IAB/GARM standards for compliance
- Use AI validation loops for efficiency

Open Questions:

- Should validation occur pre-flight, in-flight, or both?
- How can review be automated without stifling creativity?

Why it Matters:

As AI begins generating ads, validation prevents bias, misinformation, and spec violations before they scale.

Overlap:

Feeds **Brand Safety** and **Measurement**, since validated creative is core to safe delivery.
Needs **Process Oversight** for human review of exceptions.

Governance Architecture depends on a shared creative validation API and standardized rejection codes so every platform interprets creative rules in the same way.

Sub-Working Group members:

Pia Malovrh - Celtra

Alex Prokofiev - Advertible

Simon Safhalter - Celtra

Amber Hunsaker - Triton Digital

4. Process & Human Oversight* philosophy of how we build

Group Working Doc → [AdCP Human-in-the-Loop \(HITL\) Oversight Framework](#)

Definition:

Preserve human-in-the-loop oversight to prevent operational or ethical failures.

Key Themes:

- Embed checkpoints in automated workflows
- Translate human heuristics into machine-readable rules
- Define triggers that require manual review

Open Questions:

- How do we encode judgment into machine logic?
- What actions must always require human sign-off?

Why it Matters:

Even the smartest agents need circuit breakers. Oversight maintains accountability, especially where reputation or legality is at stake.

Overlap:

Underpins all domains — it dictates when humans step into **Brand Safety**, **Creative Validation**, or **Compliance** flows, and informs escalation design in **Governance Architecture**.

Governance Architecture requires a universal escalation model, triggers, autonomy tiers, reason codes, and fallback logic that all agents must use when human intervention is needed.

Sub-Working Group members:

Amy Good - Strategy

Travis Lusk - Lead

Peter Luttrell - Simpli.Fi

5. Regulatory & Compliance Alignment

Group Working Doc -  **AdCP Compliance & Regulatory Alignment Framework**

https://docs.google.com/document/d/1nmT7skA5Vonr0yTld_bQ4pwpwxeNx3jr4QN2ZsmSx1A/edit?usp=sharing

Definition:

Integrate legal, ethical, and regional frameworks directly into agentic workflows.

Key Themes:

- Build dynamic compliance libraries by vertical
- Sync with GDPR, DMA, and WFA ERA frameworks
- Automate rule updates across jurisdictions

Open Questions:

- Who maintains and authenticates compliance libraries?
 - **Match Keys** are required for interoperable systems to maintain state.
 - Whether Match Keys are subject to data protection regulations depends on whether the recipient can legally reidentify specific individuals (or households in some jurisdictions), or not. The UK ICO calls this the In-whose-hands-test, while the EEA SRB decisions and intent of Omnibus rewrite of GDPR ensures that contracts and internal company policies prohibiting (rather than “preventing” reidentification*) can ensure the information communicated among systems is not Personal Data and hence not subject to data protection regulations.
 - *Note Apple’s use of “Random Identifiers” on their consumer-facing privacy policy and omission from their CCPA disclosure of what is Personal Information. Also note FTC settlement with InMarket where the key distinction between prohibition and prevention was negotiated toward the more pragmatic option.
 - With the above in mind, having a signal about whether or not a recipient has a contract to prohibit reidentification of the match key it receives will be important scaling the cross-system, real-time communication of information.
 - **Compliance** is an obligation on each recipient to adhere to its contracts and appropriate regulations concerning the information communicated.
 - Programmatically declaring each company’s policies, similar to P3P, would aid in the facilitation of senders knowing in advance how recipients intend to collect and process information received.
 - **Enforcement** is a policing function that can only happen after violations are identified (i.e. we do not prosecute people for future crimes). As such, the automated communication layer should focus on evidence (audit trails) rather than vigilante judicial procedures.
-
- How can global rules localize without duplication?
 - Given a technical framework cannot solve conflict of law policies, I suggest that we instead focus on improving clarification of WHAT is communicated, so that various jurisdictions regulations can more easily be applied when relevant.

Why it Matters:

Automation without compliance is a liability. This ensures agentic systems respect privacy and regulation by design.

Overlap:

Pairs with **Measurement** (proving compliance) and **Manifesto** (defining permissible behavior). Works hand-in-hand with **Governance Architecture** for enforcement.

Governance Architecture needs a centralized compliance ruleset that maps jurisdictional laws to machine-readable constraints and enforces them during every agent decision.

Sub-Working Group members:

Joshua Koran - Koran Consulting - **Lead**

Eric Picard - Consultant / Compliant.global

Dinesh Bhat - HyperMindz.ai

Sean Wagner - Metatate

Lead?

6. Measurement, Verification & Validation

Group Working Doc - DRAFT  **AdCP Measurement, Verification & Validation Framework**

Definition:

Verify that campaigns ran according to declared policies and manifests.

Key Themes:

- Build real-time validation and secure enclaves
- Enable privacy-safe post-flight auditing
- Standardize compliance reporting metrics

Open Questions:

- How can verification agents confirm compliance in real time?
- Should verification be open-source or certified?

Why it Matters:

If we can't prove what happened, governance is symbolic. Verification turns promises into proof.

Overlap:

Links with **Compliance** (legal validation), **Safety** (content audits), and **Governance Architecture** (record-keeping and transparency).

Governance Architecture must define a single audit log and reporting schema that records all agent actions, rule checks, and outcomes so verification agents can prove compliance.

Sub-Working Group members:

Tina Iannacchino - Seedtag

Cem Sel - AdPriva

Peter Luttrell - [Simpli.Fi](#)

Dinesh Bhat - [HyperMindz.ai](#)

Lead?

**7. Governance Architecture & Interoperability****Definition:**

Define a scalable, neutral governance architecture spanning all AdCP components.

Key Themes:

- Operate as a meta-layer connecting creative, signals, and buy protocols
- Maintain open-source neutrality and avoid vendor capture
- Interoperate with standards like UCP and ERA

Open Questions:

- Should governance exist as its own protocol (e.g., AdGP)?
- How do we manage conflicting rules between advertisers, publishers, and regulators?

Why it Matters:

It's the backbone — the connective layer that unifies data, policy, and verification across systems.

Overlap:

Touches every domain: stores **Manifests**, enforces **Safety** and **Compliance**, connects **Measurement** outputs, and defines how the whole ecosystem talks to itself.

Not needed as a 7th area, the ideas can be folded into 1-6.

Each subgroup defines what needs to be true. Governance Architecture defines where it lives, how it is enforced, and how it is audited. One point is added to each area to cover off.

Overall Structure of Overlap:

Think of these as layers of a system:

- **Inputs:** Brand Manifesto, Creative Standards
- **Controls:** Brand Safety, Process Oversight, Compliance
- **Verification:** Measurement & Validation
- **Infrastructure:** Governance Architecture

Each depends on the others to make agentic advertising *safe, auditable, and scalable*.

Each layer reinforces the others, making agentic advertising not just automated, but *accountable, interoperable, and trustworthy*.

12/12

Terrific progress to date:

Area 1 Brand Safety: Lead [@Lauren Carbone](#) and team are working on adapting their framework in the same format as 2 & 4. Lauren mentioned discussing with Brian to pause on bringing in IAS/DV for now.

Area 2 Brand Manifesto: Lead [@Amy Good](#) and team have created the framework, [@Pia Malovrh](#) and Brian have made comments/suggestions. I encourage the remainder of the team to do the same. Brian suggested we consider this framework as an independent piece vs building it in. The sub-group will need to get together to discuss further.

Area 3 Creative Standards: [@Allie Lichtenberg](#) has offered to Lead, in partnership with [@Pia Malovrh](#) and the remaining team members. They are investing standards needed here, with companies like Boltive. Framework needs are TBD.

Area 4 Human in the Loop: [@Travis Lusk](#) / [@Amy Good](#) leads, along with [@Peter Luttrell](#). We could use 1-2 more sub-group members to contribute. Framework is open for comments/suggestions.

Area 5 Compliance: [@Joshua Koran](#) has offered to lead, thank you Joshua for all your comments. Request that others in the sub-group coordinate with Joshua to begin creating the framework.

Area 6 Measurement, Verification: Several people have offered to be in the sub-group, no lead identified yet. Suggest the sub-group meet to determine a lead and steps to creating a framework.

@Allie Lichtenberg is investigating with the larger AdCP group on who has an understanding of where each group has progressed, if there is a larger AdCP timeline and whether there is a status sheet that each group can contribute a progress update.

12/2

What's Next WIP

- Ensure each sub committee has members and a lead
- Sub committees write up and align on framework
- Share frameworks with the group and link to this document
- Present frameworks to each other?
- Share with AdCP leaders to gain alignment?
- Share frameworks with engineers to begin building
- Sub committees meet with and/or check work with engineers along the way
- Complete by when? Incremental AdCP updates? Monthly? Quarterly?

11/7 notes

[\[AG1\]](#) Need seller input. Creative agent involvement here. Goal go DTP for innovative formats. Brand define source of truth for their assets, Brand define what spots, etc to run. Brand call out if machine or human make optimizations.

[\[AG2\]](#) Define escalation requirements, may be different, more governed in the beginning vs further into the campaign

[\[AG4\]](#) Add F-caps, alignment across publishers. Can the tech be built to accommodate, will the pubs be able to implement?

Notes from 11/6 call with Amy, Travis, Lauren

Define for ourselves what the group is vs isn't – all other agents can tap into governance areas

Rules given as inputs before going live, instead of after RFP process; buyers and sellers would love it, would marketers? Do they trust the machine to keep their brand secrets?

Need interpretation cheat sheets: Marketer language, simplify; Buyer language, build confidence in tool

Be real about the upfront work for client and agency, the first time takes a while, build template/checklist

What about domains, exclusion list, AI check them

Does AdCP make us start over or evolve the process – solve and standardize the current process, evolving helps buyers adopt

Can the client/agency upload their brand manifesto, buying guidelines etc without filling it in spots in the platform

Rule of thumb: Build an audit trail

Consider governance areas are mostly in the RFP & Planning stage; 4th stage could be Analytics & Performance reporting

Manage f-caps?

From policy to practice: The three-stage governance engine

